ISCR Seminar Series

The ISCR hosted 70 seminars from visitors in FY 2004 covering a wide spectrum of research areas, and recruited an additional 35 speakers from LLNL ranks to speak to visiting students. The ASC Institute for Terascale Simulation Lecture Series was established in 2000 to enrich the intellectual atmosphere of LLNL's large simulation community through the visits of leaders representing the diverse areas of computation. In FY 2004, we hosted William Wulf, John Grosh, Bjarne Stroustrup, Mary Wheeler, and David Bailey. The ISCR Summer Student Lecture Series was also established in 2000 and forked into three different series in Summer 2004—Internships in Computational Modeling at the Terascale (ICMT), Internships in Computer Science at the Terascale (ICST), and College Cyber Defenders Computer Security. Seminar reports listed in the following table have been omitted from the printed version of the ISCR Annual Report in the interest of space, but can be found on the accompanying PDF. You can also obtain a CD-ROM containing these reports by calling (925) 423-3691.

To view the full abstract of any of the seminars listed on the following pages, simply click on the name of the seminar.

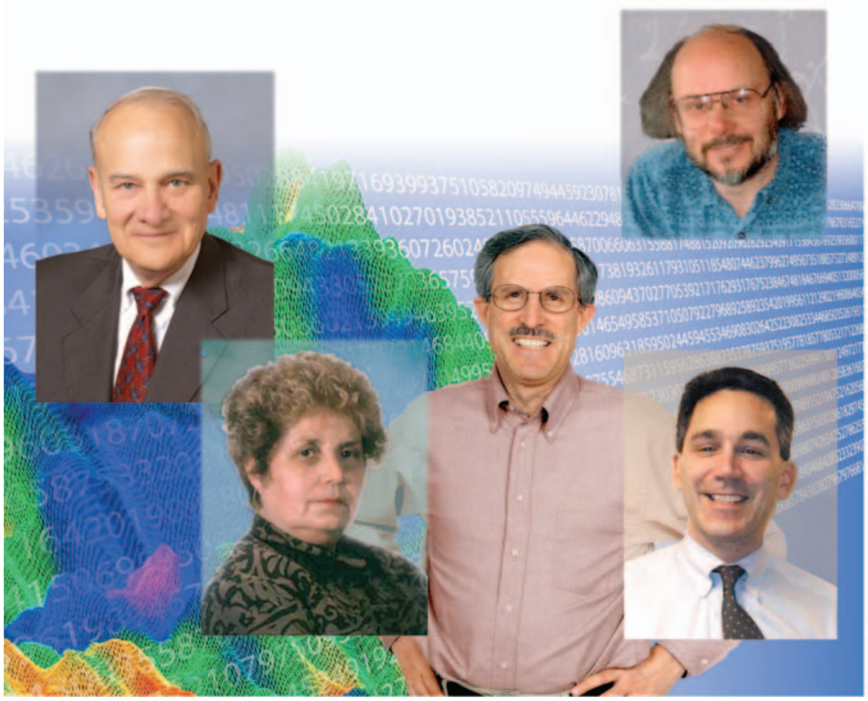

ISCR Seminar Series

Date	Speaker, Affiliation / Title of Seminar

ITS Lecture Series

Date	Speaker, Affiliation / Title of Seminar
12/4/03	**William Wulf**, National Academy of Engineering Challenges for Computing and Information Technology in the 21st Century
2/17/04	**John Grosh**, Department of Defense A DoD Perspective on High-End Computing
3/1/04	**Bjarne Stroustrup**, Texas A&M University Programming, Language and Libraries
6/24/04	**Mary Wheeler**, University of Texas at Austin Mathematical and Computational Modeling of Multiphysics Couplings
8/4/04	**David Bailey**, Lawrence Berkeley National Laboratory Twelve Ways to Fool the Masses: Back to the Future

Visitor Seminars

Date	Speaker, Affiliation / Title of Seminar
10/1/03	**Kenneth Chiu**, Indiana University Bloomington A C++ Reflection Library
10/3/03	**Branden Fitelson**, University of California, Berkeley A User-Friendly Probability Machine with Applications
10/9/03	**Olav Beckmann**, Imperial College London,United Kingdom Dynamic Code Generation in C++ as a Foundation for Domain-Specific Optimization
10/23/03	**Charles Elkan**, University of California, San Diego Progress in Clustering
10/23/03	**James Glimm**, Stony Brook University Statistical Models for Simulation Errors and Their Role in Prediction and Uncertainty Quantification
10/24/03	**Xiaolin Li**, Stony Brook University Simplification, Adaptivity and Conservation in Front Tracking Method
10/27/03	**Achi Brandt**, Weizmann Institute of Science, Israel Image Processing: From Segmentation to Recognition
10/30/03	**Douglas Swesty**, Stony Brook University Efficient Solution of the Discretized S_N Time-Dependent Boltzmann Transport Equation on Parallel Platforms
10/31/03	**Frank Mueller**, North Carolina State University Detailed Cache Coherence Characterization for OpenMP Benchmarks
11/7/03	**Daniel Segre**, Harvard Medical School Computational Models of Cellular Metabolic Fluxes

Date	Speaker, Affiliation Title of Seminar
11/13/03	**Richard Scalettar,** University of California, Davis Linear Algebra and Quantum Simulations
11/14/03	**Marsha Berger,** New York University Cartesian Grids with Embedded Geometry
11/24/03	**Johannes Gehrke,** Cornell University Techniques for Processing Large Data Streams
12/3/03	**John Feo,** Cray Inc. Cray Cascade Project
12/3/03	**Ravi Samtaney,** Princeton Plasma Physics Laboratory The Magneto–Hydrodynamic Richtmyer–Meshkov Instability
12/9/03	**Margot Gerritsen,** Stanford University Why Are Streamline Methods Attractive for Simulation of Gas-Injection Processes?
12/9/03	**Jonas Nilsson,** Stanford University A Hybrid High-Order Method for the Incompressible Navier–Stokes Equations
12/12/03	**Barney Maccabe,** University of New Mexico Early Experience in Splintering Communication Protocols
12/12/03	**Jennifer Widom,** Stanford University The Stanford Data Stream Management System
12/17/03	**Daniel Crichton,** Jet Propulsion Laboratory A Data Grid Framework for Managing Planetary Science Data
12/17/03	**William Dally and Patrick Hanrahan,** Stanford University Merrimac: Supercomputing with Streams
12/18/03	**Terran Lane,** University of New Mexico From Security to Cells: Ongoing Machine Learning Research at the University of New Mexico
1/9/04	**Christopher Jermaine,** University of Florida Approximate Query Processing with Sampling and Pre-Aggregation
1/12/04	**Thomas Seidman,** University of Maryland, Baltimore County Hybrid Systems: Discontinuous Dynamics in a Continuous World
1/16/04	**Serge Belongie,** University of California, San Diego Three Brown Mice: See How They Run — Monitoring Rodent Behavior in the Smart Vivarium
1/23/04	**Cheryl McCosh,** Rice University Type-Based Specialization in a Telescoping Compiler for MATLAB
1/26/04	**Elizabeth Post,** Lincoln University On the Farm: Parallel Small Talk for Simulating Dairy Operations

ISCR Seminar Series

Date	Speaker, Affiliation Title of Seminar
1/27/04	**Steven Parker,** University of Utah What's New with SCIRun2?
1/29/04	**Miguel Argaez,** University of Texas at El Paso An Optimization Technique for Large-Scale Nonlinear Programming
1/29/04	**Leticia Velazquez,** University of Texas at El Paso A Global Optimization Technique for Solving Zero or Very Small Residual Nonlinear Least-Squares Problems
2/10/04	**James Hobart,** Classic Systems Solutions Designing for Usability
2/10/04	**Alan Laub,** University of California, Davis Statistical Condition Estimation
2/11/04	**Demet Aksoy,** University of California, Davis PLASMA (PLAnetary Scale Monitoring Architecture)
2/18/04	**Donald Schwendeman,** Rensselaer Polytechnic Institute Numerical Method for High-Speed Reactive Flow on Overlapping Grids
2/20/04	**Matteo Pellegrini,** University of California, Los Angeles PROLINKS: A Database of Co-Evolving Proteins
2/25/04	**Laxmikant Kale,** University of Illinois, Urbana-Champaign Adaptive Resource Management via Processor Virtualization: Charm++ and AMP
3/3/04	**Lada Adamic,** Hewlett-Packard Laboratories How to Search a Social Network
3/4/04	**Christoph Pflaum,** Universität Erlangen-Nürnberg, Germany 3D Computation of Laser Cavity Eigenmodes by Finite Elements
3/5/04	**Joel Saltz,** Ohio State University Middleware Support for Data Ensemble Analysis
3/10/04	**Stephen Neuendorffer,** University of California, Berkeley Actor-Oriented Metaprogramming
3/15/04	**Luiz DeRose,** IBM Research DPOMP: An Infrastructure for Performance Monitoring of OpenMP Applications
3/18/04	**Boleslaw Szymanski,** Rensselaer Polytechnic Institute Analyzing Evolution of Virulence through Spatially-Explicit Epidemic Models
3/26/04	**Gunther Weber,** University of California, Davis Topology-Based Exploration of Scalar Fields

Date	Speaker, Affiliation Title of Seminar
4/6/04	**Kirk Hays and Max Alt,** Intel Corporation Performance Analysis and Tuning of MDCASK and PF3D Codes on Itanium Processors
5/6/04	**David Liu,** University of California, Berkeley GridDB: A Data-Centric Overlay for Scientific Grids
5/14/04	**Scott Baden & Jacob Sorenson,** University of California, San Diego Data-Driven Execution of Communication Tolerant Algorithms
5/17/04	**Homer Walker,** Worcester Polytechnic Institute Globalization Techniques for Newton–Krylov Methods
5/24/04	**Jeffrey Heys,** University of Colorado at Boulder Numerical Issues When Modeling Fluid-Elastic Interaction in 3D with First-Order System Least Squares
5/24/04	**Chad Westphal,** University of Colorado at Boulder First-Order System Least-Squares for Problems with Boundary Singularities
6/1/04	**Martin Bazant,** Massachusetts Institute of Technology Induced-Charge Electro-osmosis
6/3/04	**Marc Barthelemy,** Commissariat á l'Energie Atomique, France Structure and Modeling of Weighted Complex Networks
6/8/04	**David Jensen,** University of Massachusetts Amherst Knowledge Discovery in Networks
6/10/04	**Steven Knight,** SCons Project SCons: A Next-Generation Build Tool
6/16/04	**Wenke Lee,** Georgia Institute of Technology Worm Detection and Response: Local Strategies and Analytical Models
7/21/04	**Frank Mueller,** North Carolina State University Detailed Cache Coherence Characterization for OpenMP Benchmarks
7/23/04	**Sameer Agarwal,** University of California, San Diego On Refractive Optical Flow
7/23/04	**Gabriele Jost,** NASA Ames Research Center What Multilevel Parallel Programs Do When You Are Not Watching
7/29/04	**David Forsyth,** University of California, Berkeley Words and Pictures
7/29/04	**Ling Liu,** Georgia Institute of Technology ReFlex: Flexible and Reliable Systems Technologies for Responding to Massively Disruptive Events

ISCR Seminar Series

Date	Speaker, Affiliation Title of Seminar
7/29/04	**Nikolaos Nikiforakis,** Cambridge University, United Kingdom Emergence of Detonation in the Flowfield Induced by Richtmyer-Meshkov Instability
8/13/04	**Chris Wiggins,** Columbia University Data-Driven Approaches for Biological Networks: Inference, Organization and Analysis
8/27/04	**Alex Schweitzer,** Universität Bonn, Germany Efficient Implementation and Parallelization of Meshfree and Particle Methods: The Parallel Multilevel Partition of Unity Method
9/8/04	**E. Ann Stanley,** Los Alamos National Laboratory Using Mathematical Models to Understand the AIDS Epidemic and Guide Policy
9/9/04	**Jia Li,** University of Alabama, Huntsville Mathematical Modeling of Malaria, Early Warning System, and Transgenic Mosquitoes
9/30/04	**Reagan Moore,** San Diego Supercomputer Center Digital Libraries and Data-Intensive Computing

LLNL Summer Seminars, ICST Computer Science

6/15/04	Pat Miller, CASC FlashMob Instant Supercomputing
6/22/04	David Jefferson, CASC The Time Warp Method of Parallel Discrete Event Simulation.
6/29/04	Kim Yates, CASC BlueGene/L: The Next Generation of Scalable Supercomputer
7/13/04	Chandrika Kamath, CASC Scientific Data Mining: The Sapphire Project
7/20/04	Tom Epperly, CASC Babel Language Interoperability Tool
7/27/04	Gregory M. Pope, CADSE Why Software Quality Assurance Practices Become Evil!
8/3/04	Kim Minuzzo, NIF NIF Control System
8/10/04	Valerio Pascucci, CASC Multiresolution Computation and Presentation of Topological Structures
8/17/04	Erick Cantú-Paz, CASC Solving Problems with Evolutionary Algorithms

Date	Speaker, Affiliation Title of Seminar

LLNL Summer Seminars, ICMT Computational Science

6/9/04	David Keyes, Columbia University and ISCR A Science-Based Case for Large-Scale Simulation
6/16/04	Gary Kumfert, CASC Power Presentations
6/30/04	Peter G. Eltgroth, CASC Overview of Research Activities in the Center for Applied Scientific Computing
7/6/04	Paul Dubois, CASC Effective Architectures for Scientific Programs
7/7/04	Steven F. Ashby, CAR Computational Science at Lawrence Livermore National Laboratory
7/14/04	David Trebotich, CASC Computational Microfluidics
7/21/04	Frank Graziani, B Division Radiation Transport in 3D Random Media: Direct Numerical Simulation
7/28/04	Don Dossa, CASC Architectural Overview of BlueGene/L Supercomputer
8/11/04	Rose McCallen and Kambi Salari, CASC DOE's Effort to Reduce Truck Aerodynamic Drag — Joint Experiments and Computations Lead to Smart Design
8/18/04	Lori Freitag-Diachin, CASC Developing Interoperable Meshing and Discretization Components
8/25/05	Rob Falgout, CASC Scalable Linear Solvers: Multigrid Methods

LLNL Summer Seminars, CCD Computer Security

6/10/04	Terry Brugger, NAIC Data Mining for Network Intrusion Detection
7/1/04	Tony Bartoletti, NAIC Identifying Internet Adversaries Despite Falsified Source Addressing
7/8/04	Bill Orvis, EE-EETD The Tale of Two Rootkits — SucKit and Hacker Defenders

December 4, 2003

Challenges for Computing and Information Technology in the 21st Century

Speaker

William. A. Wulf, National Academy of Engineering, wwulf@nae.edu

Information Technology (IT), the convergence of computing and communications technologies, has had an enormous impact on all aspects of life in the developed world. It will have even more impact in both the developed and developing world as we enter the 21st century. Powered by the unprecedented and continuing advances in microelectronics and photonics, the power and capacity of our expanding information infrastructure has risen exponentially while simultaneously its cost has also fallen exponentially. At least for the foreseeable future, the exponential pace of technology improvement is likely to continue.

In this lecture, I will explore some of the non-technical, societal challenges and opportunities posed by information technology as we enter the 21st century. I will not provide answers for these challenges or guarantees that we will exploit the opportunities, but hopefully, asking some of the right questions will provoke serious thought about them.

Institution Web page:
http://www.nae.edu/

TOC

February 17, 2004

A DoD Perspective on High-End Computing

Speaker
John Grosh, Department of Defense, john.grosh@osd.mil

Over the past several years, we have witnessed a renewed interest in high-end computing. The advent of the Earth Simulator and maturation of cluster technology have spurred intense discussion at both technical and high levels of government, industry, and academia on the future of high-end computing. Frequently, such discussions suffer from oversimplification, masking the real reason why we develop and deploy high-end computing, i.e., to advance science and engineering and to achieve specific operational capability.

In this talk, I will take a Department of Defense (DoD) perspective on this field, delving into the driving factors that will define the future of high-end computing. Central to this discussion are concepts of systems engineering, which are influencing DoD strategy in this area. I will also outline DoD activities, as well as technology and policy issues related to high-end computing. In particular, I will discuss performance measurement, importance of matching applications to systems requirements, software engineering, and a new area of concern, applications software security (as opposed to network and computer security). In addition, I will discuss some of the challenges of embedded computing systems and draw some loose parallels to high-end computing. In wrapping up the discussion, I will discuss what I view are the fundamental challenges facing high-end computing from both a technical and policy perspective.

TOC

March 1, 2004

Programming, Language and Libraries

Speaker

Bjarne Stroustrup, Texas A&M University, bs@cs.tamu.edu

We say we write our code in a programming language. For real code, that's only partially true. We write our programs in an "extended language" characterized by a set of libraries supporting our general application domains and key abstractions (e.g., graphics, linear algebra, and distribution). This talk explores the relationships between language features (e.g., classes, templates, and exceptions), programming styles (e.g., object-oriented programming and generic programming), and library design in C++. All examples will be very simple to illustrate fundamental programming and performance issues. The examples will illustrate some of the design principles of C++, and based on that, I'll make a few conjectures about the likely directions for the evolution of C++ and the International Organization for Standardization (ISO) C++ standards effort.

Speaker's Web page:
http://www.research.att.com/~bs/homepage.html

Institution Web page:
http://www.tamu.edu/

TOC

June, 24, 2004

Mathematical and Computational Modeling of Multiphysics Couplings

Speaker **Mary Fanett Wheeler,** University of Texas, Austin, mfw@ices.utexas.edu

Multiphysics couplings can happen in different ways. One may have different physical processes (e.g., flow, transport, reactions) occurring within the same physical domain, or one may have different physical regimes (e.g., surface/subsurface environments, fluid/structure interactions) interacting through interfaces. We will discuss both of these types of multiphysics couplings during this presentation. Of particular interest will be the development of interpolation/projection algorithms for projecting physical quantities from one space/time grid to another, the investigation of mortar and mortar-free methods for coupling multiple physical domains, and the coupling of non-conforming and conforming finite element methods.

Speaker's Web page:
http://www.ticam.utexas.edu/~mfw/

Institution Web page:
http://www.utexas.edu/

TOC

August 4, 2004

Twelve Ways to Fool the Masses: Back to the Future

Speaker

David H. Bailey, Lawrence Berkeley National Laboratory, dhbailey@lbl.gov

In the early 1990s when highly parallel computing technology was new, researchers, as well as computer vendors, were prone to make inflated claims of the effectiveness and performance of these systems for scientific computations. Concerned by these developments, I published a tongue-in-cheek article in 1993 called "Twelve Ways to Fool the Masses When Giving Performance Results on Parallel Computers." Such warnings were not widely heeded, and as a result, the parallel computing field lost some measure of credibility, numerous parallel computer firms failed, and government agencies trimmed funding. Now the field is recovering, and new variants such as grid computing are emerging. But once again, researchers in the field are advised not to oversell this technology and to employ rigorous methodologies when analyzing and reporting results on these systems.

Speaker's Web page:
http://crd.lbl.gov/~dhbailey/

Institution Web page:
http://www.lbl.gov/

TOC

October 1, 2003

A C++ Reflection Library

Speaker

Kenneth Chiu, Indiana University, chiuk@indiana.edu

The ability to programmatically inspect a class definition can add power and flexibility to certain software systems. For example, a middleware library can use this to convert the state of an object into a sequence of bytes that can be transmitted over the network. This ability, termed reflection, relies on classes that contain metadata describing other classes. Some languages, such as Java, provide these classes as part of the language specification. However, C++ does not. In this talk, I present a C++ reflection library. I will cover both the public interface and some of the implementation specifics. I also survey some of the other existing work in this area.

Speaker's Web page:
http://www.cs.indiana.edu/~chiuk

Institution Web page:
http://www.iub.edu/

TOC

October 3, 2003

A User-Friendly Probability Machine with Applications

Speaker

Branden Fitelson, University of California, Berkeley, branden@fitelson.org

A general mechanical procedure with a user-friendly front end for reasoning about the probability calculus is presented. The procedure is then used to solve various probabilities in probability theory. Issues of computational complexity and problem size will also be discussed. All necessary technical and historical background will be provided during the talk.

Speaker's Web page:
http://fitelson.org

Research Web page:
http://fitelson.org/research.htm

Institution Web page:
http://www.berkeley.edu

TOC

October 9, 2003

Dynamic Code Generation in C++ as a Foundation for Domain-Specific Optimization

Speaker **Olav Beckmann,** Imperial College, London, ob3@doc.ic.ac.uk

The TaskGraph Library is a C++ library for dynamic code generation that combines specialization with dependence analysis and restructuring optimization. A TaskGraph represents a fragment of code that is constructed and manipulated at runtime, then compiled, dynamically linked and executed. The TaskGraph Library is implemented purely in C++ using macros and operator overloading to define a simplified, C-like sub-language that is used for initializing TaskGraphs. The internal representation used for representing generated code is SUIF-1, and the TaskGraph library implements an API for calling SUIF's analysis and restructuring passes on the generated code.

We view the TaskGraph Library as a research tool for facilitating domain-specific runtime optimizations in scientific applications. One key distinction of this approach is its combination of runtime code specialization and restructuring optimization. Sample applications include:

- Specialization of a generic image-filtering application to a specific convolution matrix
- Automatic search for optimal tile size and loop sequence for matrix multiply and dense stencil loops
- Unrolling of loops over Morton-order matrices.

We are currently planning to use the TaskGraph library as a tool for implementing and evaluating domain-specific optimization "components" for scientific programs. We are particularly interested in exploring this idea in the context of stencil loops.

Speaker's Web page:
http://hpc.doc.ic.ac.uk/~ob3/

Institution Web page:
http://www.ic.ac.uk/

TOC

October 23, 2003

Progress in Clustering

Speaker

Charles Elkan, University of California, San Diego, elkan@cs.ucsd.edu

Clustering algorithms are widely used in data mining, but basic advances are still needed in order to build genuinely robust clustering software. In this talk, I shall describe three advances that together make the well-known *k*-means algorithm far more useful. The first advance is an optimization of *k*-means that uses the triangle inequality to avoid unnecessary distance calculations. Empirically, the optimized running time has almost no dependence on *k* and the number of clusters, so hundreds of clusters can be found in gigabyte data sets in minutes. The second advance is an extension of k-means that is far better at identifying true clusterings because it does not get stuck at local minima. Finally, the third advance is an extension of *k*-means that automatically finds the appropriate number of clusters to use. This algorithm uses a novel statistical test for the hypothesis that a data subset is Gaussian, which works regardless of the dimensionality of the data.

Speaker's Web page:
http://www.cs.ucsd.edu/users/elkan

Institution Web page:
http://www.ucsd.edu/

October 23, 2003

Statistical Models for Simulation Errors and Their Role in Prediction and Uncertainty Quantification

Speaker **James Glimm,** State University of New York, Stony Brook, glimm@ams.suny.edu

Uncertainty in numerical simulations pertains to both forward simulation and to inverse problems. Probability estimates for simulation error are used to assess mismatch between approximate solutions and observations. The straightforward application of these ideas runs into a fundamental obstacle, namely an explosion of computational requirements. Determining whether an error needs a (better) comparison solution uses more resources, and determining a statistical ensemble of errors requires even more resources. To understand the parametric dependence of the statistics on the solution parameters also requires additional resources.

We present some ideas that may allow a simplification of the program, and thus allow its usefulness in practice:

1. Even deeply nonlinear problems can utilize linear models for the error statistics
2. Composition laws allow the build-up of error models for complex problems from simpler components
3. Parametric dependence of statistics can be expressed in simple regression models
4. Separate effects leading to uncertainty can be isolated and quantified individually. Methods to accomplish this step will be explained with illustration to shock physics and porous media simulations.

Our general message is that through scienfitic understanding of the errors, we will greatly reduce the required simulation resources needed for their characterization.

Speaker's Web page:
http://www.ams.sunysb.edu/~glimm/glimm.html

Research Web page:
http://www.ams.sunysb.edu/~shock/FTdoc/FTmain.html

Institution Web page:
http://www.sunysb.edu/

TOC

October 24, 2003

Simplification, Adaptivity and Conservation in Front-Tracking Method

Speaker

Xiaolin Li, State University of New York at Stony Brook, linli@ams.sunysb.edu

I will discuss three important issues in the front tracking method. First is the simplification of the geometrical handling—we use the grid-based interface reconstruction method to reduce the topological complexity. Second, we combine the tracking and adaptive methods. We developed the interoperability between two codes, FronTier and Overture. The former handles tracking while the latter handles adaptivity. We also address the conservation issue of the front tracking method through the use of the dynamical flux.

In addition, I will report on two recent developments of the front tracking method. One is the locally grid-based tracking method. This method combines the merits of both grid-free tracking and grid-based tracking. Another advance is the extension of the three-dimensional front tracking for multi-component interfaces.

Finally, I will show some success stories of the front tracking method when it is applied to important physics and engineering problems, such as the study of the Rayleigh–Taylor instability.

Speaker's Web page:
http://www.ams.sunysb.edu/~linli

Research Web page:
http://www.ams.sunysb.edu/~shock/FTdoc/FTmain.html

Institution Web page:
http://www.sunysb.edu/

TOC

October 27, 2003

Image Processing: From Segmentation to Recognition

Speaker **Achi Brandt,** Weizmann Institute of Science, achi@wisdom.weizmann.ac.il

Image segmentation is a prerequisite for higher-level processes from motion detection to object recognition and can also be used for picture denoising, compactification, miniaturization, etc. Segmentation is difficult because objects may differ from their background by any of a variety of properties that can be observed in some, but often not all scales. Our algorithm of Segmentation by Weighted Aggregation (SWA) consists of an adaptive process in which pixels are recursively aggregated into increasingly larger scale aggregates, each having progressively longer lists of coherent properties, including statistics from sub-aggregates at all finer levels. The process, motivated by the algebraic multigrid (AMG), is enhanced by top-down (coarse-to-fine) feedback of aggregation directives.

Since the computation is recursive and done mostly at coarser levels, the algorithm costs only several dozen operations per pixel and is highly parallelizable. Experimental results demonstrate a dramatically improved segmentation over current state-of-the-art methods. Moreover, the hierarchical segmentation produced in this way is in a form directly usable by learning/recognition systems since each segment emerges equipped with a vector of numbers representing textures, standardized shapes, sub-segments with their own vectors of numbers, and other identifying features. (Joint work with E. Sharon, M. Galun, Y. Gorelick and R. Basri)

Speaker's Web page:
http://www.wisdom.weizmann.ac.il/~achi/

Institution Web page:
http://www.weizmann.ac.il/

TOC

October 30, 2003

Efficient Solution of the Discretized S_N Time-Dependent Boltzmann Transport Equation on Parallel Platforms

Speaker

Douglas Swesty, State University of New York at Stony Brook,
dswesty@blackhole.ess.sunysb.edu

Implicit S_N discretizations of the time-dependent Boltzmann transport equation give rise to large, sparse, linear systems of equations that require the use of parallel architectures. Two of the most common methods for solving these sets of equations are source iteration (commonly used in nuclear engineering) and a full-matrix approach (commonly used in astrophysics). I will present results of a comparison study of the efficiency of these two approaches on a variety of 1d and 2d test problems.

Speaker's Web page:
http://www.astro.sunysb.edu/dswesty/

Institution Web page:
http://www.sunysb.edu/

TOC

October 31, 2003

Detailed Cache Coherence Characterization for OpenMP Benchmarks

Speaker

Frank Mueller, North Carolina State University, mueller@cs.ncsu.edu

Past work on studying cache coherence in shared-memory symmetric multiprocessors (SMPs) concentrated on processor simulation, including the memory hierarchy and interconnects. With the advent of hardware counters, it has become feasible to study coherence traffic based on aggregates of events. But little work has been put forward to combine the benefits of simulation methods and hardware counters. Our work closes this gap with an approach for fine-grained coherence simulation indicating opportunities for optimizations that can subsequently be confirmed by evaluating hardware counters.

The technical contributions of this work are as follows. We introduce ccSIM, a cache-coherent memory simulator fed by data traces obtained through on-the-fly dynamic binary rewriting of OpenMP benchmarks executing on a Power3 SMP node. We explore the degrees of freedom in interleaving data traces from the different processors compared to hardware performance counters. The novelty of ccSIM lies in its ability to relate coherence traffic—specifically invalidations resulting in subsequent cache misses—to data structures and to their reference locations in the source program, thereby facilitating the detection of inefficiencies.

Our experiments demonstrate that:

(a) Cache coherence traffic is simulated accurately for SPMD programming styles as its invalidation traffic closely matches the corresponding hardware performance counters

(b) We derive detailed coherence information indicating the location of invalidations in the application code

(c) We derive opportunities for optimizations from these details leading us to program transformations that result in decreased coherence misses and execution time.

Speaker's Web page:
http://moss.csc.ncsu.edu/~mueller/

Institution Web page:
http://www.ncsu.edu/

TOC

November 7, 2003

Computational Models of Cellular Metabolic Fluxes

Speaker

Daniel Segre, Harvard Medical School, dsegre@genetics.med.harvard.edu

Evolutionary adaptation can result in biological systems that are optimal for certain tasks, compatibly with their physico-chemical constraints. In the absence of a detailed knowledge of all parameters for a metabolic network, optimality can serve as a powerful constraint to produce quantitative testable predictions. Based on the premise that bacteria such as Escherichia coli have maximized their growth rate during evolution, flux balance analysis (FBA) predicts whole-cell metabolic reaction rates (fluxes) at steady state using linear programming. However, while the assumption of optimal growth rate for a wild-type bacterium is justifiable, the same argument may not be valid for artificially perturbed strains, such as genetically engineered knockouts.

I will show how one can identify such suboptimal metabolic states by assuming that the immediate response to the perturbation is a minimization of metabolic adjustment (MOMA) with respect to the wild-type flux distribution. The study of flux differences between suboptimal and optimal metabolic network states can be useful in understanding regulatory and evolutionary responses to environmental changes and ecosystem dynamics.

Speaker's Web page:
http://genepath.med.harvard.edu/~dsegre/

Institution Web page:
http://hms.harvard.edu/

TOC

November 13, 2003

Linear Algebra and Quantum Simulations

Speaker **Richard Scalettar,** University of California, Davis, rts@sherlock.ucdavis.edu

One of the most powerful methods for understanding the physics of many interacting electrons (magnetism, superconductivity, metal-insulator phase transitions, etc.) is the determinant quantum Monte Carlo algorithm. In this approach, the electronic degrees of freedom are integrated out through the introduction of an auxiliary field, leaving an expression for the partition function that consists of a very high-dimensional integral over the auxiliary field. The integrand is a product of determinants whose values depend on the field configuration. To do the integral stochastically involves computing how the determinant changes as the field configuration changes, which ends up requiring the repeated calculation of the inverse of the matrix whose determinant is the integrand.

In this talk, I will briefly review the determinant quantum Monte Carlo algorithm and then focus on the detailed structure of the matrices that arise. I will then discuss open questions of how to do the linear algebra involved more efficiently. In particular, we are looking at an alternate approach which in principle reduces the computation of the matrix M^{-1} to that of M^{-1} on a vector and thereby the scaling of the algorithm from the cube of the matrix dimension to linear. Unfortunately, the computation of M^{-1} on a vector goes horribly awry (no convergence) in practice.

The goal of this talk is to get input into the solution of these problems.

Speaker's Web page:
http://leopard.physics.ucdavis.edu/rts/

Institution Web page:
http://www.ucdavis.edu/

November 14, 2003

Cartesian Grids with Embedded Geometry

Speaker

Marsha Berger, New York University, berger@CIMS.nyu.edu

Cartesian mesh methods are one approach to solving partial differential equations (PDEs) in complex geometry. In this approach, only cells cut by the geometry require special attention. We review Cartesian mesh methods and discuss the problem of constructing stable and accurate discretizations at the cut-cells. Computational results for realistic three-dimensional aircraft will be presented.

This is joint work with Michael Aftosmis and Scott Murman of NASA Ames Research Center, Randy LeVeque at the University of Washington, and Christiane Helzel at the University of Bonn.

Speaker's Web page:
http://www.cs.nyu.edu/cs/faculty/berger/

Institution Web page:
http://www.nyu.edu/

TOC

November 24, 2003

Techniques for Processing Large Data Streams

Speaker

Johannes Gehrke, Cornell University, johannes@cs.cornell.edu

Data management techniques for data streams have gained much importance recently. I will talk about two techniques. First, I will introduce techniques for approximately answering queries over continuous data streams with limited memory. Our method relies on randomizing techniques that compute small "sketches" of a stream that can be used to provide approximate query answers with provable error guarantees. The second part of the talk will introduce techniques for efficient integration and aggregation of historical information for archiving data.

Speaker's Web page:
http://www.cs.cornell.edu/johannes/

Research Web page:
http://www.cs.cornell.edu/database/

Institution Web page:
http://www.cornell.edu/

December 3, 2003

Cray Cascade Project

Speaker

John Feo, Cray, Inc., feo@cray.com

In 2002, the US Department of Defense initiated the High Productivity Computing Systems Project to develop a next-generation computer system capable of sustaining a petaflop. Cray Incorporated was selected as one of five commercial vendors to submit an initial design. Cray, working with its partners, has proposed Cascade, a revolutionary new computer system comprised of custom processors, a next-generation interconnection network, and an active memory system.

While still in design, the system is expected to include support for both heavyweight threads that exploit high temporal locality and lightweight threads that exploit high-spatial locality. The former will execute on processors that tolerate memory latencies through a combination of multithreading, vector, and stream processing. The latter may execute in active memory systems with PIM-like characteristics that may also be multithreaded to tolerate memory latencies. The interconnection network may be a symmetric Cayley graph network capable of high bandwidth, low latency communications. Memory will be physically distributed, but shared.

A sophisticated programming environment is proposed to assist application programmers with automatically utilizing the machine's unique processing capabilities. We expect that the global shared memory and the hardware's ability to tolerate memory latencies when executing either heavyweight or lightweight threads will eliminate many of the programming challenges confronting scientific application developers today.

In this talk, I will present the design goals for Cascade and describe the architecture and programming environment as they are currently envisioned.

Institution Web page:
http://www.cray.com/

TOC

December 3, 2003

The Magneto-Hydrodynamic Richtmyer–Meshkov Instability

Speaker **Ravi Samtaney,** Princeton Plasma Physics Laboratory, samtaney@pppl.gov

In the past two decades the Richtmyer–Meshkov (RM) instability has become the subject of extensive experimental, theoretical and computational research due to its importance in technological applications such as inertial confinement fusion, as well as astrophysical phenomena such as shock interactions with interstellar clouds. In this talk, we will present recent results from nonlinear simulations of the Richtmyer–Meshkov instability in the presence of a magnetic field.

The seminar will be divided into three segments. In the first segment, we will present a primer on compressible magneto-hydrodynamics (MHD). In the second segment, we will present numerical evidence that the growth of the Richtmyer–Meshkov instability is suppressed in the presence of a magnetic field. This is due to a bifurcation that occurs during the refraction of the incident shock on the density interface. The result is that baroclinically generated vorticity is transported away from the interface to a pair of slow or intermediate magnetosonic shocks. Consequently, the density interface is devoid of vorticity and its growth and associated mixing is completely suppressed. The third segment on the talk will focus on the numerical method to obtain the aforementioned results. We will discuss the implementation of an unsplit upwinding method to solve the ideal MHD equations with adaptive mesh refinement (AMR) using the CHOMBO framework. The solenoidal property of the magnetic field is enforced using a projection method that is solved using a multigrid technique.

Research Web page:
http://w3.pppl.gov/APDEC-CEMM

Institution Web page:
http://www.pppl.gov/

TOC

December 9, 2003

Why Are Streamline Methods Attractive for Simulation of Gas-Injection Processes?

Speaker
Margot Gerritsen, Stanford University, margot.gerritsen@stanford.edu

In the SUPRI-C research group, we are interested in the design of efficient and accurate simulation tools for compositional problems, such as those occurring in gas-injection processes. The underlying method we use is the streamline method. In this talk, I would like to motivate this choice and discuss recent extensions and improvements made to the traditional streamline method that improve its accuracy and efficiency for compositional problems.

For reliable performance prediction of gas-injection processes, it is essential to properly account for the effects of heterogeneity through the use of fine grids and accurately represent component transfer between phases. Because of the required high-grid density and the costly phase behavior calculations (i.e., flashes), gas-injection processes are computationally intensive. Streamline methods akin to Euler-Lagrangian methods are attractive for these advection-dominated processes because they alleviate time-step restrictions, are inherently parallel, and easily combined with adaptive mesh refinement strategies to reduce flashing costs. In streamline methods, the pressure equation is solved on a 3D grid, while the mass balance equations for each of the components are solved along streamlines, generated from the pressure field using Darcy's law.

We show that accurate modeling of phase behavior can be achieved if high-order upwind schemes are used along streamlines. The high order of the schemes allows the use of coarser grids along streamlines, thus reducing flash calculations. Also, the longitudinal numerical diffusion is reduced, which is especially important in compositional problems because it can significantly affect the predicted displacement efficiencies. Numerical smoothing is still present as a result of mapping solution values between streamlines and pressure grid at pressure updates. To alleviate this problem, we designed a high-order mapping algorithm that greatly reduces these mapping errors. Besides these recent advances, I will also discuss our plans for the near (and far) future.

Speaker's Web page:
http://ekofisk.stanford.edu/faculty/gerritsen/

Institution Web page:
http://www.stanford.edu/

TOC

December 9, 2003

A Hybrid High-Order Method of the Incompressible Navier-Stokes Equations

Speaker **Jonas Nilsson,** Stanford University, jnilson@stanford.edu

A hybrid high-order method of the incompressible Navier-Stokes equations has been developed in an ongoing project between the Department of Scientific Computing at Uppsala University and the Department of Mechanics at Royal Institute of Technology in Stockholm. It is based on fourth-order compact finite difference approximations of Padé type in two dimensions, combined with a spectral method based on a Fourier expansion in the spanwise direction. The Padé operators allow for the use of curvilinear staggered grids. In time, we use a second order semi-implicit scheme, where the nonlinear terms are treated explicitly. In every time step, a system of linear equations is solved for the velocity and the pressure by an outer and an inner iteration. The convergence properties of the iterative method are analyzed. The order of the method is demonstrated in numerical experiments—compute the flow in a channel, the lid-driven cavity, a constricting channel and past a circular cylinder.

Institution Web page:
http://www.stanford.edu

TOC

December 12, 2003

Early Experience in Splintering Communication Protocols

Speaker

Barney Maccabe, University of New Mexico, maccabe@cs.unm.edu

Modern computing systems offer a variety of contexts for execution of code. In splintering, functionality that is traditionally centralized and executed in a single context is broken into small pieces called splinters, which are then distributed among the execution contexts in the system. The goal is to identify a splintering of the functionality so that splinters can be distributed in a way that improves overall system performance while retaining the integrity of the original implementation.

To date, we have focused our efforts on splintering related to communication protocols across the execution contexts provided by a host processor and programmable network interface card in a computational cluster. In this context, splintering is related to OS-bypass and protocol offloading. In the talk, I will describe how splintering differs from these approaches.

Speaker's Web page:
http://www.cs.unm.edu/~maccabe/

Institution Web page:
http://www.unm.edu/

TOC

December 12, 2003

The Stanford Data Stream Management System

Speaker **Jennifer Widom,** Stanford University, widom@cs.stanford.edu

This talk will describe our ongoing work developing the Stanford Stream Data Manager (STREAM), a system for executing complex continuous queries over multiple continuous data streams. The STREAM system supports a declarative query language and copes with high data rates and query workloads by providing approximate answers when resources are limited and adapting its execution strategies automatically as conditions change. We will provide an overview of the system and our research plans, highlight several specific contributions to date, and show a demo (time and logistics permitting).

Speaker's Web page:
http://www-db.stanford.edu/~widom/

Research Web page:
http://www-db.stanford.edu/stream/

Institution Web page:
http://www.stanford.edu/

December 17, 2003

A Data Grid Framework for Managing Planetary Science Data

Speaker

Daniel Crichton, Jet Propulsion Laboratory, crichton@pop.jpl.nasa.gov

As the volume of planetary science data expands seemingly without limit, the ability to share and correlate data across geographically distributed heterogeneous repositories remains a serious challenge. An architectural software framework has been developed by the Object Oriented Data Technology (OODT) task to address this challenge and has been successfully deployed in several diverse scientific domains.

Characterized by separate technology and data architectures, this framework was used by the Planetary Data System (PDS) to enable timely delivery of 2001 Mars Odyssey data to the science community as soon as the data was released from the Mars Odyssey project. As a result, OODT is now the principal infrastructure for delivery of planetary data to the scientists for all future missions and is working with the Planetary Data System to address the huge data volume increases expected for the Mars Reconnaissance Orbiter (MRO).

Based on the success deployment of OODT to support access and management of planetary science data, the Space Physics Archive Search and Extract (SPASE) effort is considering the framework to support a data search and retrieval system for the Space Physics science community. Further demonstrating its flexibility, the framework has been infused into the National Cancer Institute's Early Detection Research Network, enabling access and sharing of data critical to the development of cancer biomarkers across 10 research institutions. In 2003, OODT was selected as Runner Up for NASA Software of the Year.

Research Web page:
http://oodt.jpl.nasa.gov/oodt-site/index.html

Institution Web page:
http://www.jpl.nasa.gov

December 17, 2003

Merrimac: Supercomputing with Streams

Speaker **William Dally and Patrick Hanrahan,** Stanford University,
billd@csl.stanford.edu, hanrahan@cs.stanford.edu

Merrimac uses stream architecture and advanced interconnection networks to give an order of magnitude more performance-per-unit cost than cluster-based scientific computers built from the same technology. Organizing the computation into streams and exploiting the resulting locality using a register hierarchy enables a stream architecture to reduce the memory bandwidth required by representative applications by an order of a magnitude or more. Hence, a processing node with a fixed bandwidth (expensive) can support an order of magnitude more arithmetic units (inexpensive). This in turn allows a given level of performance to be achieved with fewer nodes (a 1-petaflops machine, for example, with just 8,192 nodes) resulting in greater reliability, and simpler system management. We sketched the design of Merrimac, a streaming scientific computer that can be scaled from a $20,000, 2-teraflops workstation to a $20 million, 2-petaflops supercomputer and present the results of some initial application experiments on this architecture.

Speaker's Web pages:
William Dally
http://csl.stanford.edu/~billd/
Patrick Hanrahan
http://www-graphics.stanford.edu/~hanrahan/

Research Web page
http://merrimac.stanford.edu/

Institution Web page:
http://www.stanford.edu/

TOC

December 18, 2003

From Security to Cells: Ongoing Machine Learning Research at the University of New Mexico

Speaker **Terran Lane,** University of New Mexico, terran@cs.unm.edu

In this talk, I will give overviews of a number of ongoing machine-learning projects that my research group is engaged in at the University of New Mexico. I will focus primarily on three projects: Bayesian methods for intrusion detection, causal activation network reconstruction for neuroimaging data, and quantitative modeling of the RNA interference process.

The intrusion detection system (IDS) project stems out of my previous research and attempts to focus on a number of questions that have not been closely examined to date in the adaptive IDS literature. Specifically, we are interested in a semi-supervised Bayesian sensor fusion approach that represents a combination of the traditionally distinct misuse and anomaly detection views of IDS. In the neuroimaging study, we are attempting to reconstruct the neurofunctional underpinnings of cocaine dependence by isolating causal activation networks from functional magnetic resonance imaging (fMRI) data. Similar to genetic regulatory network reconstruction problems, this is a high-dimension Bayesian network structure search problem and is dramatically intractable.

I will briefly review our efforts to incorporate substantial domain knowledge into the process to constrain the search and improve plausibility of the results. Finally, I will describe a new project in which we are attempting to develop a quantitative model of the recently discovered immunological process of RNA interference (a.k.a., RNA silencing, posttranscriptional gene silencing, quelling, etc.) This process is quite exciting because it offers efficient knockdown of specific genes and may represent our first tool for directly attacking viral infections, transposons, and genetic diseases. I will briefly review this process, indicate some of the open questions, and describe potential roles for computer science and machine learning in the analysis of this process.

Speaker's Web page:
http://www.cs.unm.edu/~terran/

Research Web page:
http://www.cs.unm.edu/~terran/research/publist.shtml

Institution Web page:
http://www.unm.edu/

January 9, 2004

Approximate Query Processing with Sampling and Pre-Aggregation

Speaker **Christopher Jermaine,** University of Florida, cjermain@cise.ufl.edu

Approximation is an important data management tool, particularly when the data are so numerous that computing an exact answer to a query can be prohibitively expensive. One of the simplest and most powerful approximation techniques available is random sampling. However, sampling can provide very poor approximation quality if the importance of the various data objects is not uniform.

For example, if we are trying to estimate the average net worth of all individuals in the United States, Bill Gates and Warren Buffet cause serious problems for sampling. Standard statistical techniques aimed at overcoming this problem (such as stratification and biased sampling) are often not useful in a database environment where ad-hoc queries must be supported. In this talk, I will describe an interesting database-oriented alternative for boosting the accuracy of sampling over skewed data.

Speaker's Web page:
http://www.cise.ufl.edu/~cjermain/

Institution Web page:
http://www.ufl.edu/

January 12, 2004

Hybrid Systems: Discontinuous Dynamics in a Continuous World

Speaker

Thomas Seidman, University of Maryland, Baltimore County

We consider the adaptation of ODE theory to situations in which the state space is a hybrid of continuous and discrete components so one switches abruptly from one mode to another as discrete events occur. While the talk primarily considers some general background (and open problems), one fairly concrete example is an optimal control problem involving a bacterial population that switches discontinuously between `dormant' and `active' modes, depending (hysteretically) on the concentration of a critical nutrient. One wishes to have the bacteria consume some pollutant while minimizing the amount of nutrient supplied.

Speaker's Web page:
http://www.math.umbc.edu/~seidman

Institution Web page:
http://www.umbc.edu/

January 16, 2004

Three Brown Mice: See How They Run – Monitoring Rodent Behavior in the Smart Vivarium

Speaker **Serge Belongie,** University of California, San Diego, sjb@cs.ucsd.edu

We address the problem of tracking multiple, identical, nonrigid moving targets through occlusion for purposes of rodent surveillance from a side view. Automated behavior analysis of individual mice promises to improve animal care and data collection in medical research. In our experiments, we consider the case of three brown mice that repeatedly occlude one another and have no stable trackable features. Our proposed algorithm computes and incorporates a hint of the future location of the target into layer-based affine optical flow estimation. The hint is based on the estimated correspondences between mice in different frames derived from a depth-ordering heuristic. Our approach is simple, efficient and does not require a manually constructed mouse template. We demonstrate encouraging results on a challenging test sequence containing multiple instances of severe occlusion. (This is joint work with Kristin Branson and Vincent Rabaud.)

Speaker's Web page:
http://www.cse.ucsd.edu/facresearch/facultyprofiles/BelongieS.html

Institution Web page:
http://w.ww.ucsd.edu/

TOC

January 23, 2004

Type-Based Specialization in a Telescoping Compiler for MATLAB

Speaker

Cheryl McCosh, Rice University, chom@cs.rice.edu

Telescoping languages is a strategy to automatically generate highly optimized, domain-specific libraries. The key idea is to create specialized variants of library procedures through extensive offline processing. Calls to the library in user scripts can then be replaced with calls to the appropriate variants. This talk will describe a telescoping system, called LibGen, that generates high-performance Fortran or C libraries from prototype MATLAB code.

LibGen uses variable types to guide variant generation and specialization. To solve the type-inference problem, LibGen uses a novel strategy that constructs, at each point in the MATLAB routine, a type jump-function describing the types of the local variables in terms of the types of the inputs. The type jump-functions are computed using a propositional, constraint-based formulation of the type-inference problem and an efficient, graph-theoretical algorithm to determine the solution.

This talk demonstrates the power of the approach by showing that LibGen is able to determine all the required variants for the ARPACK library, from MATLAB development code provided by the ARPACK authors. By allowing library writers to develop and maintain their code in higher-level languages, such as MATLAB, while achieving the performance of coding in lower-level languages, such as Fortran or C, telescoping languages strive to increase the productivity of the scientific community.

Institution Web page:
http://www.rice.edu/

TOC

January 26, 2004

On the Farm: Parallel Smalltalk for Simulating Dairy Operations

Speaker **Elizabeth Post,** Lincoln University, poste@lincoln.ac.nz

This talk is mainly about porting a complex system with a developed framework and components in several languages from Windows to Linux and at the same time, parallelizing it for a distributed cluster system. The core framework was implemented in Smalltalk whose virtual machine doesn't parallelize easily. This ported and parallelized model was then used for a trial optimization study using a genetic algorithm to identify the best ways of managing dairy farms. Then we had to visualize the large amounts of high-dimensional data from thousands of simulations with eight varying input parameters and one output result.

I will be talking about porting, parallelizing, using it for optimization and visualizing the results. All of this is preliminary development work—the dairy farm model is still being improved and refined to give more realistic results, although it gives some pretty reasonable results already.

Speaker's Web page:
http://www.lincoln.ac.nz/amac/profiles/poste.htm

Institution Web page:
http://www.lincoln.ac.nz/

TOC

January 27, 2004

What's New with SCIRun2?

Speaker

Steven Parker, University of Utah, sparker@cs.utah.edu

I will present an overview of a new parallel component framework (SCIRun2) and will demonstrate how software components are being used in a variety of research projects at the University of Utah. SCIRun2 is based on the Common Component Architecture (CCA) and the Scientific Computing and Imaging Institute's SCIRun and Uintah projects. SCIRun2 supports distributed computing through distributed objects. Parallel components are managed transparently over an MxN method invocation and data redistribution subsystem. A meta-component model based on CCA is used to accommodate multiple component models such as CCA, CORBA and Dataflow.

The SCI Insitute at the University of Utah utilizes component-based environments for biomedical computing, computational combustion and other applications. I will present how components are utilized in the DOE ASCI-funded Center for Simulation of Accidental Fires and Explosions (C-SAFE) to achieve scalability on thousands of processors for multi-physics simulations.

Speaker's Web page:
http://www.cs.utah.edu/~sparker/

Research Web page:
http://www.sci.utah.edu/

Institution Web page:
http://www.utah.edu/

TOC

January 29, 2004

An Optimization Technique for Large-Scale Nonlinear Programming

Speaker **Miguel Argaez,** University of Texas at El Paso, mar@math.utep.edu

We are interested in the development of efficient and robust optimization algorithms for solving large-scale nonlinear programming problems. The research plans focus on the use of interior-point methods, a path-following strategy, Krylov subspace methods, and globalization strategies. In particular, the quasicentral path notion, which is a relaxation of the central path that excludes the dual condition, is proposed as a central region to guide the iterates to an optimal solution of the problem.

For making progress towards the central region, a generalized augmented Lagrangian function is introduced as a merit function. Furthermore, a new notion of weighted neighborhoods as a measure of closeness to this central region is presented. A preliminary numerical implementation was conducted in MATLAB using the proposed methodology and is working efficiently with small- to medium-size problems, but our goal is to extend this work for solving large-scale problems.

We propose to decouple the linear system associated with the perturbed Karush-Kuhn-Tucker conditions to obtain a smaller system called a saddle point problem. We investigate efficient linear algebra solvers for obtaining an approximate solution of the reduced system that fits with the proposed interior-point Newton globalization strategy. In collaboration with a National Laboratory, we plan to develop high-quality, portable software of the proposed algorithm incorporating object-oriented software libraries. In particular, one goal is the application of the algorithm to problems of interest to the Department of Energy.

Speaker's Web page:
http://www.math.utep.edu/Faculty/mar

Institution Web page:
http://www.utep.edu/

TOC

January 29, 2004

A Global Optimization Technique for Solving Zero or Very Small Residual Nonlinear Least-Squares Problems

Speaker **Leticia Velazquez,** University of Texas at El Paso, leti@math.utep.edu

We consider global minimization of nonlinear least-squares problems with zero or very small residuals. We present an algorithm based on Levenberg-Marquardt methods combined with a multi-start technique. We compare the algorithm numerically with several existing globalization techniques, and our preliminary numerical results demonstrate that the method is quite competitive. The set of problems consists of test cases from the literature and of one well-known model problem in computational biology—the distance geometry problem. We propose to extend this approach for solving constrained nonlinear least-squares problems with nonzero residuals and to apply the algorithm to DOE data-fitting applications.

Speaker's Web page:
http://www.math.utep.edu/Faculty/leti

Institution Web page:
http://www.utep.edu/

February 10, 2004

Designing for Usability

Speaker

James Hobart, Classic Systems Solutions, Inc., jimh@classicsys.com

This presentation will explore the principles of user-centered design and how to apply these principles to your Web site or application.

Attendees will learn how to apply a proven process for identifying true user requirements, develop and validate conceptual models, and create designs that are highly usable. We will also explore the creation and usage of visual design patterns for solving complex user interface design issues.

Attendees will learn how to:

- Apply principles of user-centered design
- Use proper layout and design techniques
- Use new UI design modeling techniques
- Create and implement in-house Web standards
- Design more successful applications
- Validate and defend important design decisions
- Document design patterns to leverage what your team has learned

Speaker's Web page:
http://classicsys.com/classic_site/html/president.html

Institution Web page:
http://classicsys.com/

February 10, 2004

Statistical Condition Estimation

Speaker **Alan Laub,** University of California, Davis, laub@ucdavis.edu

Understanding the condition (or sensitivity) of problems solved with algorithms implemented in floating-point arithmetic is an essential step in assessing the accuracy of computed solutions. Standard approaches to measuring the condition of various problems in numerical linear algebra, for example, compress all sensitivity information into a single condition number. Thus, a loss of information can occur in situations where this standard condition number does not accurately reflect the actual sensitivity in a solution or in particular entries of a solution.

A method is described that overcomes these and other common deficiencies. The new procedure measures the effects on the solution of small random changes in the input data, and by properly scaling the results, it obtains condition estimates from each entry of a computed solution. This approach, which is referred to as small-sample statistical condition estimation (SCE), applies to both linear and nonlinear problems. In the former case, when an explicit Frechet derivative is available for the computed quantity in question, the method is especially efficient, costing no more than standard normwise or componentwise estimates. Moreover, SCE has the advantage of considerable flexibility. For example, it easily accommodates restrictions on or structures associated with allowable perturbations (symmetry, bandedness, etc.). The method has a rigorous statistical theory available for the probability of accuracy of the condition estimates. The theory of SCE is described along with several illustrative examples.

Speaker's Web page:
http://www.cs.ucdavis.edu/people/faculty/laub/index.html

Institution Web page:
http://www.ucdavis.edu/

TOC

February 11, 2004

PLAnetary Scale Monitoring Architecture (PLASMA)

Speaker **Demet Aksoy,** University of California, Davis, aksoy@cs.ucdavis.edu

Recent advances in wireless sensor networks enable various geo-, air- and water-based monitoring. However, monitoring events, such as disasters or environmental change, requires much more than a local deployment of a sensor network. In order to correlate the observations made by sensor networks, there is an emerging need to exchange mass amounts of globally distributed data.

Our project, PLAnetary Scale Monitoring Architecture (PLASMA), aims at developing an integrated data management and communications system for time-critical response based on real-time observations made by heterogeneous sensors. PLASMA makes use of a multi-tiered, satellite-based architecture to enable the most suitable communication for the application.

In this talk, I will present our energy-efficient and dynamic-channel allocation algorithm, which creates broadcast schedules without any initial knowledge of the network topology. Our results suggest that our online algorithm is competitive to the offline random vertex coloring algorithm that assumes perfect knowledge of the network topology while significantly improving power efficiency.

Speaker's Web page:
http://www.cs.ucdavis.edu/~aksoy/

Research Web page:
http://www.cs.ucdavis.edu/~aksoy/addtnl/publication.html

Institution Web page:
http://www.ucdavis.edu

TOC

February 18, 2004

A Numerical Method for High-Speed Reactive Flow on Overlapping Grids

Speaker
Donald Schwendeman, Rensselaer Polytechnic Institute, schwed@rpi.edu

A numerical method for the accurate simulation of high-speed reactive flow is presented. The method is an unsplit, Godunov-type, shock-capturing scheme for the reactive Euler equations, and uses overlapping grids to handle complex geometries together with a scheme of adaptive mesh refinement (AMR) to locally increase grid resolution. The method has been used to study a number of problems involving high-speed reactive flow focusing on the birth, propagation and failure of detonation waves in a variety of geometries and using a variety of reaction and equation of state models. The talk will provide an overview of the method and will discuss several calculations for a variety of reactive flow problems.

Speaker's Web page:
http://eaton.math.rpi.edu/Faculty/Schwendeman/home.html

Institution Web page:
http://www.rpi.edu/

PROLINKS: A Database of Co-Evolving Proteins

Speaker **Matteo Pellegrini,** University of California, Los Angeles, matteope@mbi.ucla.edu

Traditionally, protein evolution has been studied by measuring the sequence similarity within groups of homologous proteins. However, it is also possible to measure the co-evolution of pairs of non-homologous proteins. To infer co-evolution, we use several methods:

1. The Phylogenetic Profile technique measures the likelihood that pairs of genes are present or absent together in fully sequenced genomes
2. The Rosetta Stone technique searches for protein-fusion events
3. The gene neighbor technique looks at the intra-genic distance between pairs of genes on multiple genomes.

We will discuss statistical methods to evaluate the likelihood of co-evolution using these three approaches. We will also discuss the observation that co-evolving proteins tend to participate in common biological processes.

Institution Web page:
http://www.ucla.edu/

TOC

February 25, 2004

Adaptive Resource Management via Processor Virtualization: Charm++ and AMPI

Speaker **Laxmikant Kale,** University of Illinois at Urbana-Champaign, kale@cs.uiuc.edu

Processor virtualization and consequent message-driven execution has emerged as an extremely powerful technique in our research over the past decade. The basic idea is simple: allow the program to be expressed in terms of a large number of virtual processors and leave their assignment to physical processors on the runtime system (RTS). Yet, this simple idea empowers the runtime system to implement many adaptive strategies that enable automatic checkpointing, flexible reassignment of processors among multiple jobs, out-of-core execution, etc. Most importantly, it allows the RTS to learn the application behavior at runtime and automatically optimize performance via dynamic load balancing and communication optimizations.

These strategies have been embodied in Charm++ and Adaptive MPI (AMPI), which have been used to effectively parallelize several applications, including NAMD, a molecular dynamics program that shared the Gordon Bell prize at SC'02. I will describe Charm++/AMPI and the adaptive runtime system and present an overview of the applications being programmed with it. I will also summarize ancillary research projects based on these ideas, such as domain-specific frameworks, fault tolerance, meta-scheduling on multiple clusters with market-based bartering mechanisms.

Speaker's Web page:
http://www.cs.uiuc.edu/directory/directory.php?name=kale

Research Web page:
http://charm.cs.uiuc.edu/

Institution Web page:
http://www.uiuc.edu/

TOC

March 3, 2004

How to Search a Social Network

Speaker **Lada Adamic,** Hewlett Packard Laboratories, ladamic@hpl.hp.com

We address the question of how participants in a small-world experiment are able to find short paths in a social network using only local information about their immediate contacts. We simulate such experiments on a network of actual email contacts within an organization, as well as on a student social networking Web site.

On the email network, we find that small-world search strategies using a contact's position in physical space or in an organizational hierarchy relative to the target can effectively be used to locate most individuals. However, we find that in the online student network, where the data is incomplete and hierarchical structures are not well defined, local search strategies are less effective. We compare our findings to recent theoretical hypotheses about underlying social structure that would enable these simple search strategies to succeed and discuss the implications to social software design.

(This is joint work with Eytan Adar at HP Labs.)

Speaker's Web page:
http://www.hpl.hp.com/personal/Lada_Adamic/

Institution Web page:
http://www.hpl.hp.com/

TOC

3D Computation of Laser Cavity Eigenmodes by Finite Elements

Speaker
Christoph Pflaum, Universität Erlangen-Nürnberg, Germany,
pflaum@informatik.uni-erlangen.de

The iteration method of Fox and Li and the Gauss mode analysis are the main method for the numerical simulation of eigenmodes in laser cavities. While the Gauss mode analysis cannot be applied for the computation of complicated 3D eigenmodes, the iteration method of Fox and Li is a real 3D method. Unfortunately, this method often converges very slowly or diverges, e.g., in the case of long resonator geometries. To circumvent this problem, one would like to apply the Finite Element method directly to the Helmholtz equation (or Maxwell equation), which describes the eigenmodes in a laser cavity. The difficulty of this approach is that in the case of long resonator geometries, a large number of finite elements is needed resolve the oscillations of the eigenmodes. Furthermore, the resulting discrete equation system has a very bad condition number.

In view of these problems, we developed a new approach for the numerical approximation of eigenmodes by Finite Elements. One idea of this approach is the ansatz:

$$E(x,y,z) = \exp[i(\bar{k}-\delta)z]u_r(x,y,z) + \exp[-i(\bar{k}-\delta)z]u_l(x,y,z)$$

where $E(x, y, z)$ is the electrical field of the wave and is propagation constant of the free wave, which in the guiding structure is reduced by a small quantity. By this ansatz, the high oscillation of E are factored out, such that u_r and u_l are smooth function and can be approximated by finite elements on a relatively coarse grid. Also, u_r and u_l represent the waves traveling in the right and left direction. Another important advantage of the above ansatz is that the resulting discrete equation system can be preconditioned by relaxations in the direction of the traveling wave. Numerical results and comparisons with the Gauss mode analysis verify the accuracy of our new approach.

Institution Web page:
http://www.uni-erlangen.de/

TOC

March 5, 2004

Middleware Support for Data Ensemble Analysis

Speaker **Joel Saltz,** Ohio State University, saltz-1@medctr.osu.edu

Dramatic decreases in the cost of storage, combined with equally dramatic improvements in network connectivity will make it possible for communities to collaboratively generate and analyze very large distributed data sets. We will describe application scenarios in biomedical research, earth science and in climate modeling that motivate this work. These application scenarios will be used to provide a broad view of what advances in systems software are needed to make this vision a reality.

In many application scenarios, data sets describe sensor-acquired or simulated spatio-temporal regions. Our approach is to develop systems software able to leverage spatio-temporal descriptive metadata to support a broad range of application areas. We will describe techniques for optimized distributed data storage, indexing, retrieval and processing of ensembles of spatio-temporal data sets. We will then describe techniques for supporting on-demand data product generation and for handing spatio-temporal and relational queries directed against these data sets distributed among storage systems located in multiple parallel machines and clusters. Finally, we will describe multiple-query optimization techniques that involve grid-based semantic caching along with identification and exploitation of intermediate results shared between the queries.

Speaker's Web page:
http://bmi.osu.edu/personnel/detail.cfm?person_id=27

Research Web page:
http://www.cs.umd.edu/projects/hpsl/chaos

Institution Web page:
http://www.osu.edu/index.php

TOC

March 10, 2004

Actor-Oriented Metaprogramming

Speaker

Stephen Neuendorffer, University of California, Berkeley,
neuendor@eecs.berkeley.edu

The design of embedded software systems is significantly complicated by the system's interaction with the physical world. Concurrent interaction and the passage of time are central to this interaction. Unfortunately, most software system practices today rely on programming languages that primarily represent sequential operations that complete in an unknown amount of time. Attempts to improve the situation through operating system mechanisms, such as interrupts, timers, and threads, often introduce as many new problems as they solve. These mechanisms provide a programmer with significant flexibility, while also exposing many programming pitfalls. One approach to solving this design problem is to provide better programming languages with structured notions of time and concurrency that can be more easily understood.

I describe an actor-oriented system modeling approach that forms the basis for a better way to architect-embedded systems. An actor is a completely encapsulated component that interacts with other actors through explicit communication channels, rather than through shared memory. Because of this restriction, the interaction between actors can be understood entirely through analysis of their interfaces. The interaction between actors can also be orthogonalized from the actors themselves, allowing it to be specified separately through a Model of Computation. We specify embedded systems using Models of Computation that represent time, concurrency, and interaction with the physical world.

Although we model embedded software using behavioral components, it is important to recognize that this does not require a runtime component architecture, such as CORBA. Efficient embedded software can be generated from actor models by specializing generically specified actors to their context in a model. We call this approach to system design "Actor-oriented Metaprogramming" to distinguish it from other component-based techniques. Metaprogramming allows the use of highly generic and reusable actors for design without sacrificing implementation efficiency.

This talk will introduce Actors and Models of Computation and talk specifically about how they are implemented within Ptolemy II. I'll discuss parts of the framework that enable generic actors and give some examples of heterogeneous modeling.

Speaker's Web page:
http://www.eecs.berkeley.edu/~neuendor/

Research Web page:
http://ptolemy.eecs.berkeley.edu/

Institution Web page:
http://www.berkeley.edu/

TOC

DPOMP: An Infrastructure for Performance Monitoring of OpenMP Applications

Speaker

Luiz DeRose, IBM Research, laderose@us.ibm.com

OpenMP is today's de facto standard for shared memory parallel programming of scientific applications. However, application developers still face a large number of application performance problems, which make it harder to achieve high performance on SMP systems. These problems are difficult to detect without the help of performance tools. Unlike MPI, which has a standard monitoring interface (PMPI), OpenMP does not yet provide a standardized performance monitoring interface. In order to simplify the design and implementation of portable OpenMP performance tools, Mohr et al., proposed POMP, a standard performance-monitoring interface for OpenMP. POMP describes an API to be called by probes inserted into the application by a compiler, a pre-processor, or via a binary or dynamic instrumentation mechanism. With such a performance monitoring interface, users and tools builders can then define their own POMP-compliant libraries for performance measurement of OpenMP applications.

In this talk, I will present a POMP implementation based on dynamic probes. This implementation is built on top of DPCL, an infrastructure for binary and dynamic instrumentation from IBM. The advantage of this approach lies in its ability to modify the binary with performance instrumentation without requiring access to the source code or relinking whenever a new set of instrumentation is required. In addition, I will present two POMP-compliant libraries: Pomprof and the KOJAK POMP library, which provide the functionality for profiling and tracing OpenMP applications respectively.

Institution Web page:
http://www.ibm.com/us/

TOC

March 18, 2004

Analyzing Evolution of Virulence through Spatially-Explicit Epidemic Models

Speaker

Boleslaw Szymanski, Rensselaer Polytechnic Institute, szymansk@cs.rpi.edu

Local interactions between individual organisms influence the population dynamics of species and mediate their competition. In this talk, we describe high-performance simulation of virulence in epidemics using spatially explicit, individual-based models of multi-species habitat. We have implemented a cellular automaton model that represents individuals of two competing host species—a parasite serving as a disease vector and a vector borne pathogen. Genetic algorithms are used to model evolution with mutation and crossover operators used to simulate genetic change; we are particularly interested in the evolution of pathogen virulence. Such an approach enables us to track parentage of each organism and to account for heterogeneity of biotic and abiotic (e.g., spatial) aspects of the habitat.

We discuss the implementation of this model on parallel, distributed memory machines (e.g., IBM SP-2, or a computational cluster). The transition executed at each simulation step is complicated by the presence of genetic information in the system state, motivating the design of several parallel algorithms to cope with this model's complexity. One algorithm computes state transition probabilities efficiently based on simultaneous reduction. Another one efficiently implements the mutation process. Yet another, partitions the simulation domain among the processors.

Using the developed system, we tested a hypothesis that competition between pathogen strains drives the evolution of contagious-disease virulence. To this end, we develop a spatially detailed model of coinfection and superinfection. We assume pairwise competition between strains differing in the mortality probability that they impose on infected hosts. Coinfection happens when two or more strains of pathogen infect the same host, while superinfection refers to within-host competition of pathogens in which a more virulent strain can infect a neighboring host already infected by a less virulent strain and displace the less virulent strain from the individual. We let the probability of superinfection depend continuously on the difference in virulence between competing strains. We apply methods of adaptive dynamics to address both convergence stability and evolutionary stability.

The mean-field approximation of the spatial model predicts evolution to criticality: a strain slightly more virulent than the resident strain can always invade and exclude the resident until a further increase in virulence would imply pathogen extinction. Results of a pair approximation to the spatial model depend on the size of the neighborhood over which the pathogen is transmitted. For smaller neighborhoods, the mean-field approximation predicts a continuous evolutionarily stable set; for larger neighborhoods a single virulence is predicted that is evolutionarily and convergently stable. Simulation of the full model suggests that the pair approximation overestimates both the density of infected hosts for given virulence and the stable level of virulence.

The combined analyses indicate that increasing the size of the interaction neighborhood:

1. Increases the maximal virulence that can persist in the absence of strain competition
2. Increases the average virulence experienced by a host population (and so decreases the equilibrium frequency of infected hosts)
3. Increases the range of non-ESS virulence strains capable of ecological coexistence.

Increasing the advantage of superinfection for a given difference in virulence:

4. Increases evolutionarily stable virulence levels
5. However, it reduces the diversity of pathogen strains that may coexist.

Speaker's Web page:
http://www.cs.rpi.edu/~szymansk

Institution Web page:
http://www.cs.rpi.edu/

TOC

March 26, 2004

Topology-Based Exploration of Scalar Fields

Speaker **Gunther Weber,** University of California, Davis, ghweber@ucdavis.edu

Trivariate data is commonly visualized using isosurfaces or direct volume rendering. When exploring scalar fields by isosurface extraction, it is often difficult to choose isovalues that convey "useful" information. The significance of visualizations using direct-volume rendering depends on the choice of good transfer functions. Understanding and using isosurface topology can help in identifying "relevant" isovalues for visualization via isosurfaces and can be used to automatically generate transfer functions. Critical isovalues indicate changes in topology of an isosurface—the creation of new surface components, merging of surface components or the formation of holes in a surface component. Interesting isosurface behavior is likely to occur at and around critical isovalues.

Current approaches to detect critical isovalues are usually limited to isolated critical points. Data sets often contain regions of constant value (i.e., mesh edges, mesh faces, or entire mesh cells). We present a method that detects critical points, critical regions and corresponding critical isovalues for a scalar field defined by piecewise trilinear interpolation over a uniform rectilinear grid. We describe how to use the resulting list of critical regions/points and associated values to examine trivariate data.

Speaker's Web page:
http://graphics.cipic.ucdavis.edu/people/profile?pid=67

Institution Web page:
http://www.ucdavis.edu/

TOC

April 6, 2004

Performance Analysis and Tuning of MDCASK and PF3D Codes on Itanium Processors

Speaker

Kirk Hays and Max Alt, Intel Corporation

Intel's Itanium processor offers its users the potential to achieve high-instruction and floating-point throughput relative to clock speed through the use of ILP (Instruction Level Parallelism). Making the most efficient use of the Itanium's EPIC (Explicitly Parallel Instruction Computing) architecture, however, can be a nontrivial task for a compiler when code is ambiguous or complex.

In this seminar, two consultants from Intel show how they were able to obtain speedups ranging from 2x to more than 4x the base-optimized case with LLNL's MDCASK and PF3D codes by applying some basic optimizer pragmas and minor code changes. In nearly all cases, actual code changes are generic and can be retained on non-EPIC architectures. The optimizations presented can be directly applied to these codes' contribution to the Science Runs phase on the Thunder cluster.

The speakers will cover the following topics: project summary, runtime environment overview (including how to run and profile the codes), performance tuning process used, optimization steps, performance characteristics summary, issues encountered during optimization, code runs in parallel environment, and recommendations/ suggestions for future optimizations.

Institution Web page:
http://www.intel.com

TOC

May 6, 2004

GridDB: A Data-Centric Overlay for Scientific Grids

Speaker **David Liu,** University of California, Berkeley, dtliu@cs.berkeley.edu

We present GridDB, a data-centric overlay for scientific grid data analysis. In contrast to currently deployed process-centric middleware, GridDB manages data entities rather than processes. GridDB provides a suite of services important to data analysis—a declarative interface, type-checking, interactive query processing, and memorization. We discuss several elements of GridDB, including data model, query language, software architecture and query processing, as well as a prototype implementation. We validate GridDB by showing its modeling of real-world physics and astronomy analyses including measurements on our prototype.

Speaker's Web page:
http://www.cs.berkeley.edu/~dtliu/

Research Web page:
http://www.griphyn.org/

Institution Web page:
http://www.berkeley.edu/

May 14, 2004

Data-Driven Execution of Communication Tolerant Algorithms

Speaker
Scott Baden, University of California, San Diego, baden@cs.ucsd.edu

Many scalable applications organize into distinct phases of communication and computation, and hence, they are amenable to bulk synchronous parallelism (BSP). However, BSP is not a natural model for latency-tolerant applications that employ asynchronous communication to overlap communication with computation. In this talk, I'll present work in progress with a programming model called Tarragon that is intended to facilitate the design of latency-tolerant algorithmic formulations. I'll discuss two applications that motivate the model. The first application is an iterative kernel for finite difference elliptic solvers; the second is a Monte Carlo method for cell microphysiology.

Speaker's Web page:
http://www.cse.ucsd.edu/users/baden/

Institution Web page:
http://www.ucsd.edu/

TOC

May 17, 2004

Globalization Techniques for Newton-Krylov Methods

Speaker

Homer Walker, Worcester Polytechnic Institute, walker@wpi.edu

A Newton-Krylov method is an implementation of Newton's method in which a Krylov subspace method is used to solve approximately the linear subproblems that determine Newton steps. To enhance robustness when good initial approximate solutions are not available, these methods are usually "globalized," i.e., augmented with auxiliary procedures ("globalizations") that improve the likelihood of convergence from a poor starting point. In recent years, globalized Newton-Krylov methods have been used increasingly for the fully coupled solution of large-scale CFD problems.

In this talk, I will review several representative globalizations, discuss their properties, and report on a numerical study aimed at evaluating their relative merits on large-scale 2D and 3D problems involving the steady-state Navier-Stokes equations.

(This is joint work with John Shadid and Roger Pawlowski at Sandia National Laboratories and Joseph Simonis at WPI.)

Speaker's Web page:
http://www.wpi.edu/Academics/Depts/Math/Faculty/walker.html

Institution Web page:
http://www.wpi.edu

TOC

Numerical Issues When Modeling Fluid-Elastic Interaction in 3-D with First-Order System Least Squares

Speaker **Jeffrey Heys,** University of Colorado at Boulder, heys@colorado.edu

The mechanical interaction between a fluid and solid can be mathematically modeled using a number of different approaches depending on the physical characteristics of the problem being solved. We are interested in systems consisting of a Newtonian fluid, modeled using the Navier–Stokes equations and a linear elastic material with properties similar to a soft tissue. These coupled fluid-elastic problems are inherently nonlinear because the shape of the fluid domain is not known *a priori*, and the computational grid must be moved or mapped.

We typically use elliptic grid generation (EGG) to map the physical domain to a fixed computational domain. A FOSLS formulation of the Navier–Stokes, EGG, and linear elasticity equations provides a number of benefits to solving coupled systems problems, including optimal finite element approximation in a desirable norm (H^1), optimal multilevel solver performance, optimal scalability, and a sharp *a posteriori* error measure.

The optimality and performance of the formulation has been demonstrated extensively in 2D for a variety of problems, including the fully coupled fluid-elastic system. However, as expected, the extension to 3D brings new challenges for both the whole and the individual parts of the coupled system. Some of the issues associated with the extension to 3D have been partially or fully addressed, such as growing complexity in the multilevel solver, iteration schemes between the components of the fully coupled system, extension to parallel computers, and proper scaling of the equations. We are only beginning to answer other questions, including the handling of singularities and p-refinement.

Speaker's Web page:
http://amath.colorado.edu/faculty/heys

Institution Web page:
http://www.colorado.edu/

TOC

May 24, 2004

First-Order System Least-Squares for Problems with Boundary Singularities

Speaker **Chad Westphal,** University of Colorado at Boulder, chad.westphal@colorado.edu

Many elliptic boundary value problems have the fortunate property of a guaranteed smooth solution as long as the data and domain are smooth. However, many problems of interest are posed in nonsmooth domains and as a consequence, lose this property at the boundary. In this talk, we consider problems that have nonsmooth solutions at "irregular boundary points," that is, points that are corners of polygonal domains, locations of changing boundary condition type, or both.

Least-squares discretizations, in particular, suffer from a global loss of accuracy due to the reduced smoothness of the solution. We investigate a weighted-norm least-squares method that recovers optimal order accuracy in the weighted functional norm and weighted H^1 norm, and retains L^2 convergence even near the singularity. The method requires only *a priori* knowledge of the power of the singularity, not the actual singular solution. The theory of this general technique is studied in terms of a simplified div-curl system and shown to be similarly effective when applied to other problems. We also investigate the effect of this technique on the algebraic multigrid solvers used on the resulting linear systems.

Speaker's Web page:
http://persweb.wabash.edu/facstaff/westphac/

Institution Web page:
http://www.colorado.edu/

TOC

June 1, 2004

Induced-Charge Electro-Osmosis

Speaker

Martin Bazant, Massachusetts Institute of Technology, bazant@math.mit.edu

Induced-charge electro-osmosis (ICEO) refers to the nonlinear electrokinetic slip at a polarizable surface when an electric field acts on its own induced double-layer charge. Here, we develop a simple physical picture of ICEO in the context of some new techniques for microfluidic pumping and mixing. ICEO generalizes AC electro-osmosis at micro-electrode arrays to various dielectric and conducting structures in weak DC or AC electric fields. The basic effect produces micro-vortices to enhance mixing in microfluidic devices, while various broken symmetries—controlled potential, irregular shape, non-uniform surface properties, and field gradients—can be exploited to produce streaming flows with small AC voltages. We also present new experiments demonstrating ICEO vortices around platinum posts in polymer microchannels. Such devices may be easily integrated into biomedical microfluidics to reduce mixing times (e.g., for DNA hybridization assays).

Speaker's Web page:
http://math.mit.edu/~bazant

Institution Web page:
http://www.mit.edu/

TOC

June 3, 2004

Structure and Modeling of Weighted Complex Networks

Speaker **Marc Barthelemy,** Commissariat a l'Energie Atomique, Marc.Barthelemy@th.u-psud.fr

In addition to topological complexity, real-world networks display a gradation in the intensity strength between nodes—the weights of the links. I will present two examples— the airline connection network and the scientific collaboration network, which are representative of critical infrastructure and social system respectively. These weighted networks exhibit broad distributions and non-trivial correlations of weights that cannot be explained in terms of the underlying topological structure. These results call for the need of the modeling of complex networks, which goes beyond purely topological models. I will present a model that provides an explanation for the features observed in several real-world networks. This model of weighted network formation introduces a dynamical coupling between topology and weights by rearranging weights when a new link is introduced in the system.

Research Web page:
http://arxiv.org/find/cond-mat/1/au:+Barthelemy_M/0/1/0/all/0/1

Institution Web page:
http://www.cea.fr/default_gb.htm

June 8, 2004

Knowledge Discovery in Networks

Speaker

David Jensen, University of Massachusetts Amherst, jensen@cs.umass.edu

Networks are ubiquitous in computer science and everyday life. We live embedded in social and professional networks, we communicate through telecommunications and computer networks, and we represent information in documents connected by hyperlinks and bibliographic citations. Only recently, however, have researchers developed techniques to analyze and model data about these networks. These techniques build on work in artificial intelligence, statistics, databases, graph theory, and social network analysis, and they are profoundly expanding the phenomena that we can understand and predict.

In recent work, my students and I have discovered a variety of unique characteristics of networks, and we have developed new methods that recognize and exploit these characteristics to produce more accurate and robust statistical models. Emerging applications for these new techniques include citation analysis, Web mining, intelligence analysis, collaborative filtering, computer security, epidemiology, and financial fraud detection.

Speaker's Web page:
http://kdl.cs.umass.edu/people/jensen/

Research Web page:
http://kdl.cs.umass.edu/

Institution Web page:
http://www.umass.edu/

TOC

June 10, 2004

SCons: A Next-Generation Build Tool

Speaker **Steven Knight,** SCons Project Founder/Leader, knight@baldmt.com

A next-generation build tool, SCons, simplifies the process of building software reliably and correctly. Implemented in Python, one of the goals of the SCons project is to make software builds easy for non-programmers and programmers alike. The SCons design won the Software Carpentry build tool competition in August 2000, and the project has been attracting attention and users ever since.

The SCons project team emphasizes SCons ability to "do the right thing out of the box" while still offering enough flexibility to solve difficult build problems. This talk will contrast the SCons approach with other build tools like Make and Ant, provide an overview of SCons functionality, and discuss the project's future directions.

Speaker's Web page:
http://www.baldmt.com/~knight/ (broken)

Institution Web page:
http://scons.sourceforge.net/

TOC

June 16, 2004

Worm Detection and Response: Local Strategies and Analytical Models

Speaker

Wenke Lee, Georgia Institute of Technology, wenke@cc.gatech.edu

Worm detection systems have traditionally focused on global strategies and required a very large monitored network, say 2^{20} nodes. We consider how local networks can provide early detection and complement global monitoring strategies.

In this talk, I will first describe two local-detection approaches. The Destination Source Correlation (DSC) algorithm monitors for worm-like infection and scanning patterns to accurately detect infected hosts. The HoneyStat approach uses modified honeypots to record and analyze system and network events. The analysis result can indicate whether an automated or worm attack is present.

I will also describe our new discrete time-based worm model adapted from the AAWP (Analytical Active Worm Propagation) model. We use this model to evaluate the effectiveness of our local detection algorithms. Our analytical results show that using a /12-monitored network, a worm warning can be issued when only 0.19% of all vulnerable hosts on the Internet are infected.

We also extend our analytical model to evaluate the effectiveness of local response techniques such as traffic throttling or blocking, and patching. Our results show that with 80% deployment ratio of network level local response, Internet worms can be slowed down about five times faster than without local response. If used together with patching, worm propagation can be stopped completely.

Speaker's Web page:
http://www.cc.gatech.edu/fac/Wenke.Lee/

Institution Web page:
http://www.gatech.edu/

TOC

July 21, 2004

Detailed Cache Coherence Characterization for OpenMP Benchmarks

Speaker **Frank Mueller,** North Carolina State University, mueller@cs.ncsu.edu

Past work on studying cache coherence in shared-memory symmetric multiprocessors (SMPs) concentrates on studying aggregate events, often from an architecture point of view. However, this approach provides insufficient information about the exact sources of inefficiencies in parallel applications. For SMPs in contemporary clusters, application performance is impacted by the pattern of shared memory usage, and it becomes essential to understand coherence behavior in terms of the application program constructs, such as data structures and source code lines.

We introduce ccSIM, a cache-coherent memory simulator fed by data traces obtained through on-the-fly dynamic binary rewriting of OpenMP benchmarks executing on an SMP node. The novelty of ccSIM lies in its ability to relate coherence traffic—specifically coherence misses as well as their progenitor invalidations—to data structures and to their reference locations in the source program, thereby facilitating the detection of inefficiencies.

Our experiments demonstrate that

a. Cache coherence traffic is simulated accurately for SPMD programming styles as its invalidation traffic closely matches the corresponding hardware performance counters.
b. We derive detailed coherence information indicating the location of invalidations in the application code.
c. We illustrate opportunities for optimizations from these details.

By exploiting these unique features of ccSIM, we were able to identify and locate opportunities for program transformations, including interactions with OpenMP constructs, resulting in both significantly decreased coherence misses and savings of up to 73% in wall-clock execution time for several real-world benchmarks.

Speaker's Web page:
http://moss.csc.ncsu.edu/~mueller/

Institution Web page:
http://www.ncsu.edu/

TOC

July 23, 2004

On Refractive Optical Flow

Speaker

Sameer Agarwal, University of California, San Diego, sagarwal@cs.ucsd.edu

This paper presents a novel generalization of the optical flow equation to the case of refraction, and it describes a method for recovering the refractive structure of an object from a video sequence acquired as the background behind the refracting object moves. By structure, we mean a representation of how the object warps and attenuates (or amplifies) the light passing through it. We distinguish between the cases when the background motion is known and unknown. We show that when the motion is unknown, the refractive structure can only be estimated up to a six-parameter family of solutions without additional sources of information. Methods for solving for the refractive structure are described in both cases. The performance of the algorithm is demonstrated on real data, and the results of applying the estimated refractive structure to the task of environment matting and compositing are presented.

Speaker's Web page:
http://vision.ucsd.edu/~sagarwal

Research Web page:
http://vision.ucsd.edu/

Institution Web page:
http://www.ucsd.edu/

TOC

July 23, 2004

What Multilevel Parallel Programs Do When You Are Not Watching

Speaker **Gabriele Jost,** NASA Ames Research Center, gjost@nas.nasa.gov

The talk will compare three multilevel parallel programming paradigms suitable for shared memory computer architectures—MPI/OpenMP, MLP, and nested OpenMP. The three paradigms are applied to a set of benchmark codes from the field of Computational Fluid Dynamics. A case study will be presented using a performance analysis tool for the comparison. Detailed analysis techniques, made possible by the tool, help to differentiate between the influences of the programming model itself and other factors, such as implementation-specific operating system or architectural issues. Some implementation issues regarding the runtime support for the different programming paradigms will be discussed.

Institution Web page:
http://www.nasa.gov/centers/ames/home/index.html

Words and Pictures

Speaker

David Forsyth, University of California, Berkeley, daf@cs.berkeley.edu

It is now possible to obtain huge collections of images that carry annotations of one form or another. These annotations can take many forms. Examples include the keywords that occur in the Corel and Hemera collections using various forms of metadata—who made the artifact depicted, when it was made, etc.—that are common in museum collections; the narrative annotations that are sometimes found in museum collections; and the captions that one can collect with news images.

Such collections are interesting for two reasons. First, because visual and text information tends to be complementary, a relatively simple analysis of both the image and the text can reveal a great deal about the data item. This means that, for example, one can cluster such collections well, enabling naive users to browse a museum's collection or browse the news in a natural way. Second, such collections can be thought of as huge but poorly supervised data sets, containing both information about the appearance of objects and various forms of world knowledge. I will demonstrate a variety of methods whereby one can build probability models linking images or image regions to their annotations. With such models, one can organize a collection in a way that makes browsing easy and quite natural. One can search for pictures using words. And, what is more important, one can attach words to pictures or even to regions. Finally, one can attach names to faces.

Speaker's Web page:
http://www.cs.berkeley.edu/~daf/

Institution Web page:
http://www.berkeley.edu/

TOC

July 29, 2004

ReFlex: Flexible and Reliable Systems Technologies for Responding to Massively Disruptive Events

Speaker **Ling Liu,** Georgia Institute of Technology, lingliu@cc.gatech.edu

Emergency response and disaster recovery have always combined two elements: a planned and prepared response to regular disaster events and an ad-hoc, dynamically organized (composed) response to massively disruptive events that exceed the prepared capacity (e.g., great earthquakes, the "Perfect Storm" of 1991, 9/11 attacks). Although these large-scale natural or man-made disasters are rare events, as 9/11 showed, the consequences of such large-scale crises can be far-reaching and diverse. They cause extremely large and lasting damages. This is due partly to difficulties in effective planning, preparation, and execution of responses to mitigate damages of and recover from such enormous disasters.

Economically and politically, it is infeasible to reserve (and keep idle) large amounts of resources and infrastructure in preparation for extraordinary disasters that may or may not happen in our lifetime. Although in the past we have resigned ourselves to a fatalistic apathy towards such rare events, our reaction to the 9/11 attacks show a new resolve to face such adversity.

We propose an innovative approach that combines a low-cost preparation of potentially useful resources and infrastructure components with a dynamically composed flexible and reliable recovery infrastructure. During normal times, resources and infrastructure contribute to economically meaningful purposes. When an extraordinary disaster happens, our technologies can be used to build a flexible, effective, reliable and secure response and recovery infrastructure.

Concretely, we propose to focus on the following three technological aspects for supporting flexible and reliable risk mitigation, emergency response, and disaster recovery. First, we propose a secure overlay network infrastructure and enabling technologies to dynamically construct large-scale, self-configuring, and self-healing overlay networks for emergency response and recovery and to assist the damage containment and recovery efforts. An example of such mechanisms is trust negotiation among multiple emergency response entities. Second, we advocate mechanisms for effective risk mitigation in advance of emergencies by promoting the concept of disaster-resilient communities and techniques for building such communities. Third, but not least, we plan to develop technological solutions to support decision making and optimize the distribution of resources in the presence of inconsistent or maliciously injected conflicting information, aiming at building a flexible, reliable, and robust emergency response and risk recovery network. Flexibility is required for extraordinary disasters since we cannot predict how much damage they may cause. Reliability is required since many enormous disasters have associated secondary risks or crisis (e.g., earthquake aftershocks, correlated terrorist attacks). Robustness is required to provide resilience to coordinated attacks prepared to incapacitate or subvert the emergency response network.

In this talk, I will give a summary of the ReFlex project, a brief overview of the background knowledge on decentralized overlay networks, as well as an in-depth discussion on a selection of technical issues and some of the ongoing efforts of Reflex project at Georgia Tech.

Speaker's Web page:
http://www.cc.gatech.edu/~lingliu/

Research Web page:
http://www.cc.gatech.edu/~lingliu/#ResearchProjects

Institution Web page:
http://www.gatech.edu/

TOC

Emergence of Detonation in the Flowfield Induced by Richtmyer–Meshkov Instability

Speaker **Nikolaos Nikiforakis,** University of Cambridge, n.nikiforakis@damtp.cam.ac.uk

Combustible mixtures of gases can support two steady modes of combustion, namely, deflagration and detonation. Under certain conditions, a relatively low-speed deflagration can accelerate to form a supersonic detonation wave, a process referred to as deflagration to detonation transition (DDT). Whilst the behaviour of steady deflagrations and detonations is reasonably well understood, there are many gaps in our understanding of the nature of the transition mechanism.

The aim of this research is to investigate the transition process, i.e., the reasons behind the change of propagation mechanism from the advection/reaction/diffusion mode of a deflagration to the coupled shock/reaction system of a detonation wave and in particular the role of interfacial instabilities. To this end, the effect of the Richtmyer–Meshkov instability arising from the interaction of a shockwave with a flame has been studied by means of Implicit Large Eddy Simulations. Transition to detonation is shown to take place in the neighborhood of localised temperature perturbations (hot spots). Finally, the character of the interim combustion-driven waves arising from these hot spots is analysed.

Speaker's Web page:
http://www.damtp.cam.ac.uk/user/nn10005

Institution Web page:
http://www.cam.ac.uk/

TOC

August 13, 2004

Data-Driven Approaches for Biological Networks: Inference, Organization and Analysis

Speaker **Chris Wiggins,** Columbia University, chw2@columbia.edu

The emerging disciplines of "systems biology" or "quantitative biology" offer the promise of developing predictive and interpretable models of biological processes that could aid in the design of new biological experiments. The challenge is to extract useful information from large-scale and heterogeneous data sets in the absence of microscopic, few-parameter models.

I will give an overview of current research at Columbia using machine learning and information theoretic approaches for inferring and analyzing networks from large-scale data sets. These include "reverse-engineering" genetic networks by integrating sequence and DNA microarray data, modular organization and visualization of networks, and inference of network growth mechanisms.

Speaker's Web page:
http://www.columbia.edu/~chw2

Institution Web page:
http://www.columbia.edu/

TOC

Efficient Implementation and Parallelization of Mesh-free and Particle Methods — The Parallel Multilevel Partition of Unity Method

Speaker **Alex Schweitzer,** Universität Bonn, Germany, m.a.schweitzer@ins.uni-bonn.de

One of the reasons why mesh-free methods gained a lot of attention in recent years in both the engineering and mathematics communities is the fact that mesh-free discretizations and particle models are often better suited to cope with geometric changes of the domain of interest, e.g., free surfaces and large deformations, than classical discretization techniques such as finite differences, finite elements or finite volumes. Another obvious advantage of mesh-free discretizations is their independence from a mesh. Mesh generation is still the most time-consuming part of any mesh-based numerical simulation. However, the cost associated with numerical integration in mesh-free Galerkin methods is usually larger than for mesh-based methods.

There exists a number of variants of mesh-free methods—smoothed-particle hydrodynamics (SPH), reproducing kernel-particle methods (RKPM), element-free Galerkin methods (EFGM), radial-basis functions (RBF), partition of unity methods (PUM), and many more. All these approaches do not depend (at least to a great extent) on a mesh or any fixed relation between the discretization points or particles. The PUM is a very flexible approach due to the product structure of its shape functions. It is a general framework rather than a specific method. The assumptions on the various (independent) components involved in the construction of a PUM function space are abstract properties only, so that we can select problem-dependent components. The PUM approach allows not only for an h-version, p-version and hp-version discretization but also for the exploitation of *a priori* knowledge about the solution in the design of an optimal approximation space.

Some of the key issues involved with mesh-free Galerkin discretization techniques are the fast construction of the shape functions, the numerical integration problem, the treatment of essential boundary conditions, and the efficient solution of the arising linear systems. We will consider these issues in the context of the PUM; however, the presented concepts are applicable to most mesh-free Galerkin approaches.

Institution Web page:
http://www.uni-bonn.de/index_en.shtml

TOC

September 8, 2004

Using Mathematical Models to Understand the AIDS Epidemic and Guide Policy

Speaker **E. Ann Stanley,** Los Alamos National Laboratory, spidermoon_nm@yahoo.com

Mathematical models have been used to study the spread of many diseases. These models have provided basic knowledge of epidemic spread and insight into the design of vaccination programs and other control measures. But perhaps the impact of models on policy and research has been greatest in the case of the AIDS epidemic due to its devastating nature and complexity.

I will discuss basic epidemiological theory and then present some insights obtained from AIDS models. These insights have affected study designs for sexual behavior surveys, needle and condom programs, and other aspects of research and public policy. I will finish with a brief discussion of our most recent work on contact tracing, which is one of the most hotly debated methods for controlling HIV spread.

Institution Web page:
http://www.lanl.gov

TOC

September 9, 2004

Mathematical Modeling of Malaria, Early Warning System, and Transgenic Mosquitoes

Speaker
Jia Li, University of Alabama at Huntsville, jli@t7.lanl.gov, li@math.uah.edu

In this talk, I will briefly talk about the life cycle of malaria, present some basic mathematical malaria models, discuss the effects of environmental changes on the transmission dynamics of malaria, the impact of releasing genetically altered mosquitoes in preventing malaria and their preliminary mathematical modeling and analysis.

Speaker's Web page:
http://ultra.math.uah.edu/~li/

Institution Web page:
http://www.uah.edu/

TOC

September 30, 2004

Digital Libraries and Data-Intensive Computing

Speaker

Reagan Moore, San Diego Supercomputer Center, moore@pop.sdsc.edu

Scientific data collections that represent the digital holdings of a research community are now being assembled into digital libraries. Scientists use the digital libraries to support browsing of registered material, discovery of relevant digital entities, and display of the data. This is similar to traditional services provided by digital libraries for image and document collections. However, scientists also need the ability to support manipulation of entire collections as part of data intensive computing. Entire collections are accessed for analysis, streamed through a processing pipeline, and the results are registered back into the digital library. The additional capabilities required by digital libraries to enable data intensive computing are examined for analysis of scientific data collections.

Speaker's Web page:
http://www.sdsc.edu/~moore/rmoore.html

Institution Web page:
http://www.sdsc.edu/

June 16, 2004

FlashMob Instant Supercomputing

Speaker

Pat Miller, CASC

It used to be that building a supercomputer of any sort was a multi-million dollar undertaking only available at government laboratories, large corporations, and large academic institutions. Cluster computing in which large numbers of commodity nodes are glued together with a commodity interconnect seems, initially, to have really changed the landscape of parallel computing. However, access to parallel computing is still quite restricted because even owning a small cluster is prohibitive due to several factors, including

1. Cost
2. Experienced personnel to administer it
3. Space
4. HVAC, etc.

What do you do if you want a personal supercomputer? The answer, of course, is a temporary supercomputer in which nodes are volunteered on a temporary basis to form a large-scale, tightly-coupled computing system for the duration of program execution.

We call this new model "FlashMob Computing." A FlashMob supercomputer is constructed from the same building blocks as a conventional cluster—computational nodes and an interconnect. The key differences are that the nodes need not be uniform, as they will be donated by interested parties. The interconnect will be cobbled together from available resources, and there will be limited time to tune the constructed system. Also, in contrast to traditional supercomputing, a FlashMob system is built around a specific application so that non-experts can run significant scientific software on self-built FlashMob clusters.

We'll talk about the "live booting" Linux distro used by FlashMob I and do a live demo of the software. We'll talk about the Supercomputer-in-a-gym event held at the University of San Francisco last April and how the successful packaging of the high-performance LINPACK application has made it possible for relatively unskilled computer enthusiasts to build and benchmark small clusters.

June 22, 2004

The Time Warp Method of Parallel Discrete Event Simulation

Speaker

David Jefferson, CASC

Time Warp is one of the best known and most general algorithms for scalable parallel discrete event simulation. The simulation model is first decomposed into parallel processes, which then schedule events for one another by sending messages, each timestamped with the simulation time of the event it describes. The fundamental synchronization requirement—the problem that Time Warp solves—is that the simulation as a whole must be executed in such a way that all event messages appear to be processed sequentially, in nondecreasing timestamp order, even though they are neither sent nor arrive in sorted timestamp order, and it is not known statically which processes will send event messages to which others.

Time Warp is a very unusual algorithm in that it relies on a very general form of process rollback for synchronization instead of process blocking. At first glance, general process rollback in a distributed, message-passing environment might seem impossible, or at least extremely difficult, to implement efficiently. However, Time Warp solves the problem in an extremely elegant way using the notion of antimessages—if two antimessages are enqueued in the same queue, they both disappear and the queue gets shorter.

Time Warp was originally implemented at the Jet Propulsion Lab in the early 90s on the Mark III Hypercube and the BBN Butterfly. Performance results on benchmark problems running on up to 112 nodes will be presented. Even now, variations of Time Warp are the best methods known for general parallel discrete simulations without special structure.

TOC

June 29, 2004

BlueGene/L: The Next-Generation of Scalable Supercomputer

Speaker

Kim Yates, CASC

We are giving an overview of the BlueGene/L Supercomputer. This is a jointly funded research partnership between IBM and the Lawrence Livermore National Laboratory as part of the ASC Advanced Architecture Research Program. This massively parallel system of 65,536 nodes is based on a new architecture that exploits system-on-a-chip technology to deliver target peak processing power of 360 teraFLOPS (trillion floating-point operations per second). The machine is scheduled to be operational in early 2005, at price/performance and power consumption/performance targets unobtainable with conventional architectures.

July 13, 2004

Scientific Data Mining: The Sapphire Project

Speaker

Chandrika Kamath, CASC

The Sapphire project is developing scalable algorithms and software for the interactive exploration of large, complex, multi-dimensional scientific data. We are using ideas from data mining to improve the way in which scientists extract useful information from data. Our work focuses on research in algorithms, incorporation of this research into software, and the application of the software to real-world problems at LLNL. The needs of these applications drive our research. In this talk, I will describe what is involved in data mining.

I will discuss our work using examples from applications, such as detection of human settlements in satellite imagery, detecting and tracking moving objects in video, and finding similar objects in data.

July 20, 2004

Babel Language Interoperability Tool

Speaker
Tom Epperly, CASC

Babel is a state-of-the-art language interoperability tool for high-performance computing being developed at Lawrence Livermore National Laboratory. Babel provides efficient, bidirectional communication between programs in C, C++, Fortran 77, Fortran 90, Java and Python. Compared with industrial alternatives, such as CORBA and COM, Babel's key distinguishing features are high performance, support for Fortran, and support for scientific data types, such as complex numbers and large multi-dimensional arrays.

Babel provides a uniform object model and exception handling across all its supported languages. It uses its own Scientific Interface Definition Language to define inter-language programming interfaces. Internally, Babel achieves interoperability through the use of the intermediate object representation (IOR) in C.

By having a common intermediate representation, Babel converts the many-to-many language interoperability problem into two many-to-one problems. Object-oriented polymorphism is achieved using virtual function table implemented in the IOR. Babel is a key architectural element of the Common Component Architecture—a component model for high-performance computing.

TOC

July 27, 2004

Why Software Quality Assurance Practices Become Evil!

Speaker

Gregory M. Pope, CADSE

This paper looks at the challenge of determining the best practices for software development and why the topic usually sparks a lively debate. The premise is that best practices are application specific and are not easily portable from one industry to another. The paper considers a case study of three different types of software developments and contrasts their differences. The paper presents an alternative method to best practices, which is a common set of principles that are turned into appropriate best practices based on project risk.

TOC

NIF Control System

Speaker

Kim Minuzzo, NIF

The National Ignition Facility (NIF) at LLNL supports the U.S. Department of Energy's National Nuclear Security Administration (NNSA) Defense Programs and LLNL missions of ensuring that the nation's nuclear weapons remain safe, secure, and reliable without nuclear testing. NIF is a 192-beam laser for creating conditions of extreme temperatures and pressures in the laboratory, and when completed, it will be the largest laser in the world.

At LLNL, we are developing the distributed control system required to operate NIF, called the Integrated Computer Control System (ICCS). ICCS uses a distributed object-oriented architecture based on CORBA to provide the interface between the software components distributed across the embedded processors and the supervisory workstations used by the operators to control NIF. User interfaces are coded in Java and most of the control software is written in Ada. The VxWorks real-time operating system is used on the embedded processors. The speaker will present an overview of NIF, ICCS and the SQA practices used on the ICCS.

August 10, 2004

Multiresolution Computation and Presentation of Topological Structures

Speaker
Valerio Pascucci, CASC

Scalar fields are used to represent data in a wide variety of applications like scientific computing, geographic information systems or medical imaging. Some fundamental topological features of a scalar field are represented in a simple graph structure that is know in different fields as a Contour Tree or Reeb graph.

In this talk, I will present the first technique that allows efficient computing of an augmented version of the Contour Tree and encompasses a complete topological characterization of the 3D scalar field. The new information provided allows the user to know, for example, the number of tunnels and voids of an isosurface without rendering or even computing the isosurface itself. The complexity analysis shows how the new approach allows the computing of the added information with minimal penalty.

The most recent extensions of this work include the decomposition of the Contour Tree in a multi-resolution data structure and its progressive presentation in real time.

I will present specific examples highlighting how an interface based on the progressive contour tree allows the user to understand the "structure" of a scalar field more accurately and minimizes the time spent exploring the data interactively.

TOC

August 17, 2004

Solving Problems with Evolutionary Algorithms

Speaker

Erick Cantú-Paz, CASC

Evolutionary algorithms (EAs) are randomized search methods inspired by evolution and genetics. EAs have been used successfully in many applications for science, engineering, and business. In this talk, I will introduce the basic evolutionary algorithms and present several applications of EAs to data-analysis problems at LLNL. I will also discuss other interesting applications of EAs in art and engineering design.

TOC

June 9, 2004

A Science-Based Case for Large-Scale Simulation

Speaker **David E. Keyes,** Columbia University and ISCR

The July 2003 report "A Science-based Case for Large-scale Simulation" (SCaLeS) from the U.S. Department of Energy (DOE) documents new levels of importance for supercomputing in scientific discovery. Together with four other federal reports on supercomputing within the past year, it made the case for an ultrascale computational facility as the DOE's second-highest priority facilities project over the next 20 years, following only the International Thermonuclear Experimental Reactor (ITER, www.iter.org) in importance.

The SCaLeS report takes a look at DOE's Scientific Discovery through Advanced Computing (SciDAC) program—a collection of 51 interconnected projects throughout science, mathematics, and computer science—halfway through its five-year scope, evaluates its early success in mixing the "cultures" of science applications (e.g., astrophysics) with enabling technologies (e.g., adaptive meshing tools), extrapolates the potential of simulation in the eyes of the applications scientists, and outlines hurdles to realizing that potential. It concludes with eight major recommendations. Approximately 315 computational scientists, applied mathematicians, and computer scientists contributed to SCaLeS.

In this talk, the editor of the SCaLeS report summarizes its findings and recommendations, and highlights some of the computational science and software infrastructure of the SciDAC initiative. This software targets multi-teraflop/s platforms and much of it is freely available. Scalable solution algorithms for simulations based on systems of partial-differential equations are emphasized, and are one of the recognized challenges for overall application scalability and of special interest to the speaker.

SCaLeS Web site:
http://www.pnl.gov/scales

Project Web site:
http://www.tops-scidac.org

Speaker Web site:
http://www.columbia.edu/~kd2112/

June 16, 2004

Power Presentations

Speaker

Gary Kumfert, CASC

For this talk, think of me as Dear Abbey. Good presentation skills are like good manners. It's something you have to be taught. It's something you have to practice consistently. It's not for your benefit, but for the benefit of people you interact with. It's most effective when it appears natural and effortless. And it says something about either how well educated you are and/or how highly you regard your audience.

Professional researchers are professional presenters. Presentations are a major tool in building your reputation, acquiring funding and defending your work when funding is up for review. This talk will cover slides, props (e.g., podium, pointers), showmanship, and lecture hall etiquette. It's a good refresher even for experienced speakers, and it's good practice for me since students are quick to point out when I don't follow my own advice!

June 30, 2004

Overview of Research Activities in the Center for Applied Scientific Computing

Speaker **Peter G. Eltgroth,** Director, CASC

An overview of the Center for Applied Scientific Computing (CASC) is presented. It includes a brief description of CASC activities in advanced software technology, applied mathematics, computational physics, computer science, data science, numerical methods, scalable algorithms for parallel processing, and scientific computing. Some technical detail is provided for ongoing projects in multigrid methods for linear system solution, code transformations, graph theory, adaptive mesh refinement, and computational biology.

July 6, 2004

Effective Architectures for Scientific Programming

Speaker

Paul Dubois, CASC

Scientific programming is dominated by change. The challenge is to develop scientific simulations in a way that allows them to remain correct despite a high and persistent rate of change not faced by programmers in other disciplines. This talk discusses a variety of ideas related to this challenge.

TOC

July 7, 2004

Computational Science at Lawrence Livermore National Laboratory

Speaker **Steven F. Ashby,** Director, Computing Applications and Research

Large-scale simulation plays an increasingly important role in many scientific and engineering applications, including those in the defense, energy, and life sciences. The time and length scales of interest in many physical and biological processes span several orders of magnitude, thus requiring the use of sophisticated numerical methods and massively parallel computers.

Dr. Ashby will discuss some of the challenges inherent in terascale simulation with an emphasis on the computational science activities at Lawrence Livermore National Laboratory (LLNL), which currently houses more scientific computing horsepower than any other facility in the world. After giving a brief overview of the Laboratory, he will describe the LLNL simulation environment and highlight some of its computer science and mathematics research efforts. The majority of the presentation will be spent surveying a number of examples of computational science at the leading edge—showing how large-scale scientific simulation is being used to advance scientific discovery. In particular, this presentation will emphasize the importance of combining simulation, theory, and experimentation via illustrative case studies drawn from a variety of scientific and engineering application areas.

July 14, 2004

Computational Microfluidics

Speaker

David Trebotich, CASC

Microfluidics is the next-generation technology for miniaturization, portability, and networking of the macroscale fluidic systems currently used in "chem/bio" applications. The underlying physics—complex biological flow at the microscale—is not well understood. Also, design of microdevices is lengthy and costly due to imperfections in the micro-fabrication cycle. The on-demand ability to predict the behavior of biological macromolecules in microenvironments will give significant advantage over current trial-and-error techniques and available rudimentary design tools. The computational microfluidics effort at LLNL will help improve design times and costs, as well as provide an optimization and advanced design platform by developing new algorithms to model complex fluids at the microscale and integrating those algorithms into a design tool.

TOC

July 21, 2004

Radiation Transport in 3D Random Media; Direct Numerical Simulation

Speaker
Frank R. Graziani, B Division

The transport of radiation through materials characterized by large-scale heterogeneities is important in a wide range of applications. In high-energy density physics, it is the propagation of photons through turbulent media. In atmospheric physics, it is solar radiation coupling to the ocean and land mass through clouds. In astrophysics, it is the transport of photons from young stars through turbulent molecular clouds.

In this talk, we will briefly discuss the tremendous theoretical effort that has gone into characterizing this problem. We will show that despite this effort, many problems of relevance to applications are still unsolved. Using the computer as a laboratory, we show how direct numerical simulations of transport through random media allow one to understand and characterize the transport behavior. In particular, using optical-path distribution functions calculated numerically allows one to compute effective mean-free paths and their fluctuations. This tool allows insight into theoretical models and offers hope in understanding the transport of photons in random media in the laboratory.

TOC

July 28, 2004

Architectural Overview of BlueGene/L Supercomputer

Speaker

Don Dossa, CASC

An overview of the hardware and software architecture of LLNL's next supercomputer will be discussed. This 360-teraflop system features several new approaches to massively parallel computing. In addition to describing the system, several areas of applications will also be discussed to show how this system will increase the level of sophistication and breadth of simulations at LLNL.

TOC

August 11, 2004

DOE's Effort to Reduce Truck Aerodynamic Drag — Joint Experiments and Computations Lead to Smart Design

Speaker **Rose McCallen and Kambiz Salari,** CASC

At 70mph, overcoming aerodynamic drag represents about 65% of the total energy expenditure for a typical heavy truck vehicle. The goal of our U.S. Department of Energy supported consortium is to establish a clear understanding of the drag-producing flow phenomena. This is being accomplished through joint experiments and computations, leading to the "smart" design of drag-reducing devices. This presentation will describe our objective and approach, provide an overview of our efforts and accomplishments, and discuss our direction for the future.

August 18, 2004

Developing Interoperable Meshing and Discretization Components

Speaker

Lori Freitag-Diachin, CASC

Typically, the first step in numerically solving a PDE-based application is to generate a discrete representation of the computational domain (the mesh) and approximate the continuous differential operators and solution field on that mesh (the discretization). Over the years, a variety of mesh generation tools and discretization methods have been developed that offer different advantages and disadvantages for different application regimes. In addition, advanced techniques and software tools that provide adaptive mesh refinement, time-varying meshes, mesh-to-mesh data transfer, and parallel decomposition of the mesh have been shown to have significant impact on the application areas that employ them. In each case, the software tools providing these advanced capabilities are becoming increasingly accepted by the scientific community, but it is often not clear *a priori* which techniques are best suited to solve a particular application problem. Ideally, the application scientist should be able to easily insert and experiment with a number of different meshing and discretization software tools, but the application programming interfaces are rarely compatible, making experimentation a labor-intensive and error-prone code modification process.

The Terascale Simulation Tools and Technologies (TSTT) Center has been funded by the Department of Energy's SciDAC initiative to address the barriers preventing easy interoperability and interchangeability of multiple mesh and discretization strategies within a single simulation. We are focusing our effort on the creation of common interfaces for existing TSTT Center technologies that will provide better interoperability and fundamentally increase capabilities that allow application scientists to easily switch among them. I describe the current status of our interface definition effort, the tradeoffs required to balance performance and flexibility, the tools used to address language interoperability issues, and our approach to simplifying the adoption process. To ensure the relevance of our research and software developments, we collaborate closely with both SciDAC application researchers and other technology centers. In particular, I will describe the use of the TSTT interfaces and philosophy to insert advanced adaptive mesh refinement (AMR) and error-estimation procedures into an accelerator-modeling code. This code will develop a new capability that combines front tracking and AMR to deploy a TSTT-compliant mesh quality improvement toolkit into three mesh generation codes in as many days.

TOC

August 25, 2004

Scalable Linear Solvers: Multigrid Methods

Speaker **Rob Falgout,** CASC

In this talk, we discuss multigrid methods development at LLNL. Multigrid methods are scalable primarily because they have optimal work requirements. That is, a linear system with N unknowns can be solved by a multigrid method with O(N) work. This property makes it possible to solve ever larger problems on proportionally larger parallel machines in constant time. Classical iterative methods like conjugate gradients are not scalable.

We discuss geometric multigrid methods for solving structure-grid problems, algebraic multigrid methods (AMG) for unstructured-grid problems and parallelization issues for both. The majority of the talk will focus on AMG, a multigrid approach that already works remarkably well over a wide variety of applications. In some cases, however, AMG is not effective without making certain problem-specific modifications and carefully tuning parameters. To address this, CASC researchers have developed new AMG algorithms, new parallelization techniques, and new theory to guide future algorithm development.

TOC

June 10, 2004

Data Mining Methods for Network Intrusion Detection

Speaker

Terry Brugger, NAIC

Network intrusion detection systems have become a standard component in security infrastructures. Unfortunately, current systems are poor at detecting novel attacks without an unacceptable level of false alarms. We propose that the solution to this problem is the application of an ensemble of data mining techniques that can be applied to network connection data in an offline environment, augmenting existing real-time sensors.

This talk will expand on this motivation, particularly with regard to operating in an offline environment and our interest in multisensor and multimethod correlation. We'll then review existing systems, from commercial systems to research-based intrusion detection systems. Next, we'll survey state-of-the-art methods in the area. Standard data sets and feature extraction turned out to be more important than we had initially anticipated, so each will be covered in some depth. Next, we will review the actual data mining methods that have been proposed or implemented. We'll conclude by summarizing the open problems in this area, along with some questions of a broader scope. We hope that by providing the motivation and summarizing the work in this area we can stimulate further research.

Terry Brugger is a Ph.D. student at the University of California, Davis, where he's researching Network Intrusion Detection Systems in addition to his day job as a computer scientist for Lawrence Livermore National Laboratory's Information, Operations and Assurance Center. Terry is the principal investigator for LLNL's College Cyber Defenders (CCD) program, which gives college students real-life experience in computer security technologies.

TOC

July 1, 2004

Internet Ballistics: Identifying Internet Adversaries Despite Falsified Source Addressing

Speaker **Tony Bartoletti,** NAIC

Network scanners constantly probe the Livermore network looking for vulnerabilities. Analyzing packet arrival timing data reveals highly distinctive patterns that may correlate with the attacker's choice of tools, physical platform and/or network location. Consistent identification will improve network security and aid counterintelligence efforts. We have developed tools to preprocess scan data using wavelet techniques to achieve a more than 1,000x compression ratio while still preserving essential features. Initial experiments indicate our methods consistently identify patterns in the data.

July 8, 2004

A Tale of Two Rootkits — SucKit and Hacker Defender

Speaker

Bill Orvis, EE-EETD

Two different rootkits have been found in compromised systems during the last year; SucKit (SK) and Hacker Defender (HkDef). SK is a Linux rootkit and HkDef is a Windows rootkit. Both of these rootkits implement new and interesting methods of hiding themselves and the intruder's files and processes. They also implement hidden backdoors that would not be detected by normal scanning. This paper describes the operation of the two rootkits, how they hide files and processes, and how they implement a hidden backdoor. It also discusses how to detect these rootkits and how to scan for the hidden backdoor.

TOC